

AD-A281 466



11

Discourse Interpretation  
and the Scope of Operators

DTIC  
ELECTE  
JUL 13 1994  
S F D

Massimo Poesio

Technical Report 518  
June 1994

This document has been approved  
for public release and sale; its  
distribution is unlimited.

1548



9421298

UNIVERSITY OF  
ROCHESTER  
COMPUTER SCIENCE

94 7 12 05 8

# Discourse Interpretation and the Scope of Operators

by

Massimo Poesio

Submitted in Partial Fulfillment

of the

Requirements for the Degree

DOCTOR OF PHILOSOPHY

Supervised by

Professor Lenhart K. Schubert

Department of Computer Science  
College of Arts and Science

University of Rochester  
Rochester, New York

1994

Accession For		
NTIS	CRA&I	<input checked="" type="checkbox"/>
DTIC	1AB	<input type="checkbox"/>
Unannounced		<input type="checkbox"/>
Justification .....		
By .....		
Distribution /		
Availability Codes		
Dist	Avail and / or Special	
A-1		

This material is based upon work supported by the National Science Foundation under Grant CDA-8822724, by ONR/DARPA Grant N00014-92-J-1512, and by the U.S. Air Force Rome Laboratory Contract F30602-91-C-0010. The government has certain rights in this material.

# REPORT DOCUMENTATION PAGE

Form Approved  
OMB No. 0704-0188

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.

1. AGENCY USE ONLY (Leave blank)		2. REPORT DATE June 1994	3. REPORT TYPE AND DATES COVERED technical report	
4. TITLE AND SUBTITLE Discourse Interpretation and the Scope of Operators			5. FUNDING NUMBERS ONR/DARPA N00014-92-J-1512	
6. AUTHOR(S) Massimo Poesio				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Computer Science Dept. University of Rochester 734 Computer Studies Bldg. Rochester, NY 14627-0226			8. PERFORMING ORGANIZATION REPORT NUMBER TR 518	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) Office of Naval Research Information Systems Arlington, VA 22217 ARPA 3701 N Fairfax Drive Arlington, VA 22203			10. SPONSORING/MONITORING AGENCY REPORT NUMBER	
11. SUPPLEMENTARY NOTES				
12a. DISTRIBUTION/AVAILABILITY STATEMENT Distribution of this document is unlimited.			12b. DISTRIBUTION CODE	
13. ABSTRACT (Maximum 200 words) (see title page)				
14. SUBJECT TERMS natural language processing; ambiguity; scope; context-dependency; discourse processes; semantic interpretation; models; task- oriented conversation; quantification; definite descriptions			15. NUMBER OF PAGES 144	
			16. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT Unclassified	18. SECURITY CLASSIFICATION OF THIS PAGE Unclassified	19. SECURITY CLASSIFICATION OF ABSTRACT Unclassified	20. LIMITATION OF ABSTRACT U	

## Curriculum Vitae

Massimo Poesio [REDACTED] attended the University of Torino, Italy from 1981 to 1986, and graduated *summa cum laude* with a Diploma di Laurea in Informatica (Bachelor of Science Degree in Computer Science) with a dissertation on knowledge representation techniques for natural language processing systems under the supervision of Prof. P. Torasso.

The author then spent one year working on speech understanding at CSELT, the Research Laboratory of the Italian National Telephone Company, and two years at the University of Hamburg, Germany, participating to the discourse understanding project WISBER under the direction of Dr. H. Marburger, Prof. B. Neumann, and Prof. W. Wahlster.

He came to the University of Rochester in the Fall of 1988 as a PhD student in Computer Science, pursuing his research under the direction of Prof. Lenhart K. Schubert and participating to the TRAINS project under the joint direction of Prof. Schubert and Prof. James F. Allen.



## Acknowledgments

I started working on the topic of this dissertation my very first semester in Rochester. I had a couple of ideas on scope ambiguity I wanted to try out, so I talked with Len Schubert about them, and we decided I could start working on that while looking for a 'real' dissertation topic. In the next five years, Len listened to a lot of weird ideas, always patiently explaining why he thought these ideas wouldn't work, until finally I arrived at a theory that resembles a lot what he had been suggesting from the very beginning.

James Allen provided the right contrast to Len. While Len always wants a theory that ties in with everything else we know about natural language, James always pushes for the clear, simple idea that can be immediately understood. I keep having a hard time in doing that, but if I have gotten any better, it's because of James.

The other members of my committee played an important role as well. Alessandro Zucchi got me interested in formal semantics, taught it to me, and always pushed me to raise my standards. Mike Taxenhaus convinced me that often people do not process natural language as someone trained in natural language processing and logic would expect them to do. Peter Laenzlein could always find the time for a discussion. And Jeff Pelletier kept asking questions that looked very simple but were very hard to answer.

A number of people outside my committee helped me shape my ideas about ambiguity, scope, and discourse. I particularly wish to thank Howard Kurtzman, Robin Croper, Graeme Hirst, Herbert Clark, Uwe Reyle, Dan Hardt, Chris Barker, David Dowty, Megumi Kameyama, Ricky Passonneau, David Milward, and Mitch Marcus. I am grateful to Ron Brachman and Henry Kautz for their support, enthusiasm, and for what they taught me about research in one summer at Bell Labs. I would also like to add that being able to rely on such detailed analyses of semantics and discourse such as those due to Hwang and Schubert, on the one side, and to Kamp and Reyle, on the other, made my work much easier.

I am grateful to both James and Len for building in Rochester a very attractive environment for work in natural language, planning and knowledge representation. The TRAINS project has grown into an extremely exciting framework to do research in natural language, and all of us working on TRAINS have become better and better at working like a team. I wish to thank David Traum and Chung Hye Hwang, that routinely find counterexamples in my most solid theories; George Ferguson for all the discussions about planning and monotonic reasoning, and all the help with the hickey stick and the frisher; Marc Light, for his endless enthusiasm for linguistics discussions; and Peter Ilicman, who can make sense of whatever input my module produces.

The natural language community in Rochester does not include only us in Computer Science, but also faculty and students in Linguistic, Philosophy, and Psychology; and, thanks to Tim

Hever and Beverly, Derek, Graham and Roberto, our activities are not limited to attending lectures and seminars. Thanks to all for all the work and all the parties.

That I ended up really enjoying my stay in Rochester has a lot to do with my discussions with Roberto, dancing and dining with Polly, running with Virginia, hiking with Lambert, and moviegoing and partying with all of them. There always was a party at Nat Martin's, Jim Muller's, or the Rarnfeld. The Little Theatre and the Dryden helped, too.

I completed this dissertation during my stay in Berkeley, for which I have to thank Jerry Feldman, and my current stay in Edinburgh, where I got invited by Robin Croper. The environment in Edinburgh is very special, and the nights are short, all of which made for a very productive stay.

Finally, I wish to thank my mom and dad, who have been calling long distance every week since I left Italy, even when I didn't return their calls; Alberto, who would wait for me when cross-country skiing and teach me origami; and Alessandro and Elena, even though I can't see Maria as often as I wished. Thanks also to that other outcast from Torino, Barbara, with whom I may have exchanged more than a thousand email messages now.

But most of all, I wish to thank Janet for choosing Rochester over Yale.

## Abstract

The problem of ambiguity is central in any theory of language interpretation, whether we are interested in language processing in humans or in developing a viable natural language processing system. Psycholinguistic evidence suggests that human subjects are able to choose an interpretation when necessary, and that competing factors are involved in this choice; however, no theory of language interpretation deals satisfactorily with the *combinatorial explosion paradox*—the fact that no matter how ambiguous natural language sentences are, they are usually interpreted without significant effort.

The main idea presented in this dissertation is that the scope preferences observed in the literature are not the result of an independent 'scope disambiguation' module, but of independent interpretation processes such as definite description interpretation or the interpretation of modals. None of these interpretive procedures is especially concerned with 'scope disambiguation', but the result of these inferences is that relations of *contextual dependency* such as anaphoric reference or presuppositionality become part of the common ground, the scope preferences observed in the literature reflect these relations of dependency. The dissertation includes a formal proposal concerning the representation of contextual dependency, and its impact on the semantics of sentence constituents.

My theory of ambiguity is based on a distinction between *semantic ambiguity*, that can be captured implicitly, by means of *underspecified representations*, and *perceived ambiguity*, that results from the process of discourse interpretation. My model of the common ground can be used to characterize both situations characterized by the presence of semantic ambiguity, and situations characterized by the existence of perceived ambiguity.

The reasoning that leads to the establishment of scopng preferences makes use, I argue, of information that is pragmatic in nature; this calls for a model of discourse interpretation in which the common ground contains such information. In the case of spoken language conversations the common ground must be a model of the *discourse situation* of the conversational participants

## Table of Contents

Curriculum Vitae	ii
Acknowledgements	iii
Abstract	v
List of Tables	xiii
List of Figures	xiv
1 Introduction and Outline of the Theory	1
1.1 AMBIGUITY IN NATURAL LANGUAGE	4
1.1.1 Kinds of Ambiguity	4
1.1.2 Ambiguity, Vagueness, and The Combinatorial Explosion Paradox	5
1.1.3 Ambiguity and Perceived Ambiguity	8
1.2 AVOIDING THE COMBINATORIAL EXPLOSION PARADOX	8
1.2.1 Syntactic and Semantic Constraints	8
1.2.2 Is There Such a Thing as 'Scopal Ambiguity'?	11
1.2.3 Underspecified Representations	13
1.2.4 Beyond the Underspecification Hypothesis	16
1.3 TOWARDS A THEORY OF SCOPE DISAMBIGUATION: DISAMBIGUATION FACTORS	18
1.3.1 Linear Order	18
1.3.2 Grammatical Function, Quantifier Hierarchy and Topic	18
1.3.3 Scope and C-Command	20
1.3.4 Van Lehn's Results	21
1.3.5 The Use of Disambiguation Factors in NLP Systems	22
1.3.6 Kurtzman and MacDonald's Experiments	23
1.3.7 Summary	26

1.4	SCOPE DISAMBIGUATION IN CONTEXT: THE TRAINS DIALOGUES	26
1.4.1	The TRAINS Corpus	27
1.4.2	Operators in the TRAINS Corpus	28
1.4.3	The Scope of Definite Descriptions	31
1.4.4	The Scope of Modals	31
1.4.5	Syntactic and Semantic Constraints	32
1.4.6	Further Observations	32
1.5	OVERVIEW OF THE PROPOSAL	34
1.5.1	Recapitulation	34
1.5.2	Discourse Interpretation as Inference over Underspecified Representations and the DRS Construction Algorithm	34
1.5.3	Contextual Interpretation and Scope	37
1.5.4	The Results of Karttman and MacDonald, Revisited	41
1.5.5	Interactions Between the Discourse Interpretation Procedures That Affect Scope and the Architecture of Discourse Interpretation	44
1.5.6	Hypothesis Filtering and the Role of Commonsense Knowledge in Discourse Interpretation	46
1.5.7	Reflexive and Abductive Architectures for Discourse Interpretation	50
1.5.8	SAD 93	50
1.6	CONTENT OF THE DISSERTATION	51
2	Formal Tools for Discourse Interpretation, Or: From Discourse Representation Theory and Situation Theory to Conversation Representation Theory	52
2.1	A CRASH COURSE ON DISCOURSE REPRESENTATION THEORY	53
2.1.1	'Donkey' Sentences and Discourse Representation Structures	53
2.1.2	The DRT Construction Algorithm	55
2.1.3	Inference in DRT	60
2.2	SITUATION THEORY	62
2.2.1	Situations and Situation Inclusion	62
2.2.2	Meaning and Information	62
2.2.3	Situation Types	63
2.2.4	Partiality	63
2.2.5	Situation Semantics	64
2.2.6	Parameters	65
2.3	EPISODIC LOGIC	65
2.3.1	The Language of Episodic Logic	65
2.3.2	Truth at a Situation in Episodic Logic	66
2.3.3	Additional Syntactic Constructs	67
2.3.4	From English to Episodic Logic	67

3	Conversation Representation Theory	69
3.1	FROM DRT TO CRT	69
3.2	THE INPUT TO DISCOURSE INTERPRETATION: LOGICAL FORMS	70
3.2.1	Lexical Semantics and the Input to Discourse Interpretation	76
3.2.2	Adding Logical Forms to DRT	73
3.2.3	Semantic Ambiguity and The Semantics of Logical Forms	75
3.3	DESCRIBING PHRASE STRUCTURE	80
3.4	CONTEXT-DEPENDENCY AND PARAMETERS	81
3.5	EVENTS, SITUATIONS AND DISCOURSE REPRESENTATION STRUCTURES	83
3.5.1	Situation Descriptions	84
3.5.2	Events and Event Descriptions	84
3.5.3	Situation Kinds	85
3.6	PRESUPPOSITION AND THE INTERPRETATION OF NOUN PHRASES	86
3.6.1	Strong and Weak NPs	86
3.6.2	Presuppositions	87
3.6.3	Presupposition as Anaphora	88
3.6.4	Presupposition, Parameters and Situation Descriptions	89
3.6.5	Existential Presuppositions and Strong Noun Phrases	91
3.7	DISCOURSE INTERPRETATION AND THE PRAGMATICS OF DISCOURSE	93
3.7.1	Definite Description Interpretation and the Common Ground	93
3.7.2	The Common Ground as the Representation of a Discourse Situation	95
3.7.3	Focus of Attention and Visual Attention	96
3.7.4	Common Ground and Mental States	97
3.8	DISCOURSE SEGMENTATION AND CONVERSATIONAL EVENTS	97
3.8.1	Discourse Segmentation and Speech Acts	98
3.8.2	Speech Acts in Natural Language Processing	98
3.8.3	Speech Acts in Situation Theory and Conversation Representation Theory	99
3.8.4	Conversational Threads, Discourse Topics, and Discourse Segmentation	100
3.8.5	Conversational Event Generation Rules	102
3.8.6	Micro Conversational Events	106
3.9	WEAK AND STRONG DEFAULTS IN DISCOURSE INTERPRETATION	111
3.10	SUMMARY: A REVISED MODEL OF DISCOURSE INTERPRETATION	112

## 4 Formal Aspects of CRU

4.1 TYPES AND THE MODEL	114
4.1.1 Types in CPT	114
4.1.2 Objects and Their Structure	116
4.1.3 Situations	116
4.1.4 The Model	117
4.2 DENOTATION AND ENTAILMENT	118
4.3 DEFINITION OF THE EVALUATION FUNCTION	119
4.3.1 Variables, Non-Logical Constants, Discourse Markers, and Parameters	120
4.3.2 Application	121
4.3.3 Equality	122
4.3.4 Connectives	123
4.3.5 Lambda Abstraction	125
4.3.6 DRSS	128
4.3.7 Situation Descriptions	131
4.3.8 Quantifiers	135
4.3.9 Conditionals and Adverbs of Quantification	138
4.3.10 Modals	140
4.4 DRS CONSTRUCTION RULES AS INFERENCE RULES	141
4.4.1 Construction Rules as Operations on DRSS	141
4.4.2 Weak and Strong Inference Rules	144
4.4.3 Syntactic and Semantic Definitions of Inference	145

## 5 Syntactic Structure and Lexical Semantics in a TRAINS fragment

5.1 NOUN PHRASES THAT ARE NOT OPERATORS: PROPER NAMES, PRONOUNS, AND BARE PLURALS	146
5.1.1 Proper Names	146
5.1.2 Pronouns	147
5.1.3 Bare Plurals	148
5.1.4 Logical Forms for Non-operator Noun Phrases	148
5.2 EVERY-NPS AND OTHER QUANTIFIERS	149
5.2.1 The Semantic Translation of Strong and Weak Determiners	149
5.2.2 Logical Forms	149
5.3 DEFINITE DESCRIPTIONS	151
5.3.1 The Semantics of Definite Descriptions: Two Proposals	151
5.3.2 Interpretation of Definite Descriptions: The Location Theory	153

5.3.3 A Situation-Theoretic Formulation of the Semantics of Definite Descriptions	155
5.3.4 Logical Forms	155
5.3.5 Plural Definite NPs	155
5.4 INDEFINITE NOUN PHRASES	157
5.4.1 Presuppositional and Non-Presuppositional Translations for Indefinite NPs	157
5.4.2 Logical Forms	158
5.5 VERBS AND VERB PHRASES	159
5.5.1 Thematic Roles	159
5.5.2 Semantic Translation of Predicates	162
5.5.3 Syntax of Verb Phrases and Logical Forms	162
5.6 TENSE	163
5.6.1 The Syntax of Tense	163
5.6.2 Contextual Factors in Tense Interpretation and the Semantics of Tense	164
5.6.3 The Proposed Semantics for Tense	167
5.6.4 Logical Forms	168
5.7 ADVERBIALS	169
5.7.1 The Syntax of Adverbials	169
5.7.2 The Semantics of Adverbials	170
5.7.3 Logical Forms	171
5.8 INFINITIVAL CLAUSES	171
5.8.1 The Syntax of Infinitival Clauses	171
5.8.2 The Semantics of Infinitival Clauses	171
5.9 MODAL VERBS AND MODAL AUXILIARIES	171
5.9.1 The Syntax of Modal Auxiliaries	172
5.9.2 The Syntax of Modal Verbs	178
5.9.3 The Semantics of Modal Auxiliaries and Modal Verbs	178
5.9.4 Logical Forms	182
5.10 THERE-INSERTION SENTENCES	183
5.10.1 The Syntax of There-Insertion Sentences	184
5.10.2 The Semantics of There-Insertion Sentences and the Definiteness Restriction	185
5.10.3 The Definiteness Restriction and Scope	186
5.11 THE MEANING OF SENTENCES	187
5.11.1 The Cooper Value of CN	187

5.11.2	A Revision	190
5.12	OTHER CONSTRUCTIONS AND SENTENCE CONSTITUENTS	191
5.12.1	Negation	191
5.12.2	Coordination	192
5.12.3	Apposition	193
5.13	SUMMARY OF THE CHAPTER	194
6	Discourse Interpretation and Scope Disambiguation	195
6.1	THE DISCOURSE SITUATION IN THE TRAINS CONVERSATIONS	197
6.1.1	(Visual) Information about the Map	198
6.1.2	The Information about the Plan	199
6.1.3	Discourse Topics: Situations vs. Situation Types	200
6.1.4	Summary	201
6.2	THE STATUS OF INFORMATION IN DISCOURSE: THE GIVEN/NEW/AR TICULATION, AND SCOPE DISAMBIGUATION	202
6.2.1	The Given/New Contract	202
6.2.2	Formalizing the Given/New Contract	203
6.2.3	Givenness and Scope Preferences	204
6.2.4	The Mapping Hypothesis	205
6.3	ARGUMENT SELECTION, THE THEMATIC HIERARCHY, AND THE EVENT DESCRIPTION BUILDING PROCEDURE	206
6.3.1	Argument Selection and the Thematic Hierarchy	207
6.3.2	An Event Description Building Procedure	208
6.3.3	Scope and the Thematic Hierarchy	211
6.3.4	Additional Predictions	212
6.4	THE INTERPRETATION OF OPERATORS	213
6.5	DEFINITE DESCRIPTION INTERPRETATION	213
6.5.1	Model Construction Rule	214
6.5.2	The Interpretation of 'Visible Situation' Definite Descriptions	215
6.5.3	Anaphoric Uses of $\infty$ -finite Descriptions	220
6.5.4	Definite Description Interpretation and Scope Disambiguation	221
6.6	THE INTERPRETATION OF MODAL AUXILIARIES AND MODAL VERBS	225
6.6.1	Model Construction Rules	225
6.6.2	Discourse Interpretation Rules for Modals	227
6.6.3	Plans as Discourse Topics	230
6.6.4	Modal Subordination	231

6.6.5	Modal Interpretation and Scope Preferences	234
6.7	TENSE INTERPRETATION	235
6.7.1	The Model Construction Rule for Tense	235
6.7.2	Identification of the Course of Action	236
6.7.3	Tense Interpretation and Scope Disambiguation	238
6.8	OTHER QUANTIFIERS	240
6.9	INDEFINITES	240
6.9.1	Model Construction Rules	240
6.9.2	Specificity and the Scope of Indefinites	241
6.10	ADVERBIALS	242
6.10.1	Model Construction Rules for Adverbials	242
7	Implementation: The Scope and Deindexing Module of TRAINS-93	245
7.1	A GENERAL OVERVIEW OF SAD-93	246
7.1.1	The Role of SAD-93 within TRAINS-93	246
7.1.2	SAD-93 Input and Output	247
7.1.3	The Architecture of SAD-93	250
7.1.4	The Dialogues Handled by the System	252
7.2	THE SYNTAX OF SAD-EXPRESSIONS	254
7.3	DETAILS ABOUT THE IMPLEMENTATION	258
7.3.1	Hypothesis Generation	260
7.3.2	Rules	266
7.3.3	The Fact Database	266
7.3.4	Discourse Markers and Priming	268
8	Conclusions	269
	Bibliography	270

## List of Tables

1.1	Determiners found in the TRAINS transcripts, June 6, 1991	29
1.2	Connectives found in the TRAINS transcripts, June 6, 1991	30
1.3	Tense, auxiliaries, and negation in the TRAINS transcripts, June 6, 1991	30
1.4	Wh phrases in the TRAINS transcripts, June 6, 1991	31

## List of Figures

1.1	The map used by the participants in the conversation	27
1.2	An unedited TRAINS transcript	29
1.3	One view of discourse interpretation	45
1.4	The interaction of strong and weak defaults	47
6.1	The common ground in the TRAINS conversations---a summary	201
7.1	The architecture of TRAINS-93	246
7.2	The Architecture of SAD-93	251
7.3	The dialog processed by TRAINS-90, TRAINS-91 and TRAINS-92	252
7.4	The dialog processed by TRAINS-93	253
7.5	The map used by the participants in the conversation	254
7.6	BNF Definition of the Syntax of SAD-expressions (1)	255
7.7	BNF Definition of the Syntax of SAD-expressions (2)	256
7.8	BNF Definition of the Syntax of SAD-expressions (3)	257
7.9	Generation of new hypothesis trees	258
7.10	Pruning hypothesis trees	260
7.11	Hypothesis propagation algorithms	261
7.12	The model construction rule for present tense	263
7.13	The model construction rule for universals	265
7.14	The rule application algorithm	267

## b Why don't you ask for help?

The listener (or reader) of an ambiguous sentence is faced with the task of recovering the interpretation(s) intended by the speaker or writer of the sentence. Ambiguous sentences may have a huge number of alternative interpretations, yet humans have no problems in processing them: this is known as the *combinatorial explosion paradox*.

The ambiguity of natural language sentences raises problems of interest for linguists, psychologists and developers of natural language processing systems alike. Which sentences are in fact ambiguous, and in which way? On which occasions is the intended interpretation of an ambiguous utterance in fact recovered by the listener/reader? How is this intended interpretation arrived at?

Although a large literature exists on the subject of ambiguity, these questions are still largely unanswered. In fact, some have argued that these questions are not even worth looking into. It is a common misconception that sentences like (1.1), (1.2a) and (1.2b) are the exception rather than the norm, and therefore ambiguity is an issue that we should tackle after other, more central questions about natural language understanding have been solved. Even a superficial perusal of any corpus of natural language texts will, however, reveal that this idea is misguided. Innocent looking but nevertheless ambiguous sentences like (1.1), are exceedingly common in such corpora. Some naturally occurring examples are shown below, all taken by a corpus of transcripts of naturally occurring conversations, the TRAINS corpus collected at the University of Rochester [Cross *et al.*, 1993]. In sentence (1.3a), the pronoun may be anaphoric either on "an engine" or on "the boxcar". That sentence is also structurally, lexically and scopally ambiguous (see below). Sentences like (1.3b) are often used in the TRAINS conversations, whose participants share the same knowledge about the world, to suggest a possible way to accomplish a certain task (i.e., in answer to questions like "How can we get the oranges to Bath?"); but in conversations like those of the Map Task corpus collected at the University of Edinburgh [Thompson *et al.*, 1993], where the participants to a conversation usually have different information, similar sentences can also be used to give new information to the other participant.

- (1.3) a. We should hook up an engine to the boxcar in Dansville, and move it to Bath.  
b. There is a boxcar at Corning.

While it is true that ambiguity is usually resolved in context (as these examples show), we cannot base our theory of ambiguity on the assumption that that is always the case. Most linguistic puns are based on our ability to perceive ambiguity. Abbott and Costello's routine "Who's on first," for example, is only funny for a listener/reader who can detect the ambiguity of the sentences in the routine. Ambiguity can also be exploited for rhetorical effects: the sentence "Can we?", the punch line of a campaign of the Edinburgh City Council promoting recycling of tin cans, deliberately plays on the ambiguity of the lexical item "can".

Several theories of surface discourse interpretation are now available [Hobbs, 1979; Hirsch and Martin, 1987; Charniak and Goldman, 1988; Hobbs *et al.*, 1990; Pereira and Pollack, 1991; Dalrymple *et al.*, 1991; Alshawi, 1992; Hwang and Schubert, 1993; Kamp and Reyle, 1993]. The developers of these theories have looked in detail at discourse processes such as

<sup>1</sup>A sentence may be deliberately ambiguous: see below.

## 1 Introduction and Outline of the Theory

Suppose Mr. Rice, addressing Mr. Porter, utters sentence (1.1).

- (1.1) I can't find a piece of paper.

The intention of Mr. Rice in using this sentence may have been to inform Mr. Porter that he is looking for any piece of paper on which to write, say, the instructions for Mr. Porter so that he can get to their common friend Mr. Richardson's house, or that he is looking for a specific piece of paper—maybe the piece of paper on which he had written Mr. Richardson's phone number. Different actions may be required of Mr. Porter depending on the actual intention of Mr. Rice. If the utterance is felicitous (i.e., no miscommunication occurs), Mr. Porter will infer the interpretation of (1.1) intended by Mr. Rice. A central question for natural language processing and psycholinguistics is: How (and when) does Mr. Porter recognize Mr. Rice's intention?

In recent years, a lot has been learned about some aspects of the interpretation process especially about the extraction of sentence constituents (parsing) and the contextual interpretation of lexical items. We still lack, however, a clear picture of a number of processes involved in what I'll call *surface discourse interpretation*—the "process by which humans use what they heard, the context, and their knowledge about language and the world to arrive at the intended interpretation of an utterance, or to realize that there has been a miscommunication."

One of the most important lessons learned by those working on surface discourse interpretation is that most natural language sentences<sup>1</sup> are ambiguous in more than one way: a speaker may mean different things when using them, and the appropriate interpretation is usually only recoverable in context. (1.1) is not exceptional in this sense. (1.2a) can be used either to describe a situation<sup>2</sup> in which the ref. part of the pronoun "he" lowers herself or a situation in which some people see the pet water/fowl of a female person. (1.2b) may be interpreted either as a question or as a suggestion.

- (1.2) a. They saw her duck.  
b. They saw her duck.

<sup>1</sup>The term *sentence* is sometimes used in the generative literature in a technical sense to indicate a tuple of disambiguated levels of representation: each tuple represents a completely disambiguated analysis of a certain sentence string. In that framework, ambiguity is not a property of sentences but of sentence strings. In this dissertation, however, I am going to use the term *sentence* in its intuitive meaning of string of words, and I talk about *surface discourse ambiguity*.

<sup>2</sup>I am using here the term *situation* in an informal sense.

reference resolution and tense interpretation, and have addressed important issues such as the syntax/semantics interface, or the role of commonsense knowledge in discourse interpretation. Many of these theories do not, however, address the problem of ambiguity, and those that do (e.g., [Charniak and Goldman, 1988, Hobbs *et al.*, 1990]) concentrate on providing tools for using whatever information one has to arrive at a disambiguated interpretation. None of the questions about the disambiguation process I mentioned above are looked at in any detail, nor do these accounts include a study of the factors that play a role in disambiguation—instead, they explain how, given a list of such factors and their relative importance, a single interpretation may be arrived at. This is particularly true for one kind of disambiguation process, the process of assigning a scope to operators—logical expressions such as quantifiers, modals, negation, or adverbs of quantification.

In this dissertation I am concerned with two tasks. First of all, I develop a theory of (surface) discourse interpretation meant to address the question that none of the existing theories of discourse interpretation answers satisfactorily: how do humans deal with ambiguous sentences? To this end, I introduce a distinction between *semantic ambiguity* and *perceived ambiguity*, and propose an account of the *combinatorial explosion paradox* mentioned above—the fact that ambiguous sentences may have a huge number of alternative interpretations, yet humans have no problems in processing them.<sup>4</sup>

Secondly, I use this model of discourse interpretation to develop an account of *scope disambiguation*. Scope disambiguation is the surface discourse interpretation process involved in the interpretation of utterances that contain more than one operator, and is arguably the least understood aspect of surface discourse interpretation. I concentrate in particular on *scope* ambiguity in context, and study the relation between scope disambiguation and other discourse interpretation processes in detail.

Many of the issues discussed in this dissertation arose in the context of the TRAINS project at the University of Rochester, that studies issues of language comprehension, planning, and reasoning encountered in task oriented natural language conversations [Allen and Schubert, 1991, Traum *et al.*, 1994]. In this dissertation I also describe a computer implementation of my proposals about disambiguation that is used as a component of the TRAINS-93 discourse understanding system.

In the remaining sections of this chapter I will discuss the problem of ambiguity in more detail. I will examine the existing proposals about scope disambiguation, and present an overview of my approach to surface discourse interpretation, ambiguity and scope disambiguation.

Although this work was originally motivated by a problem discussed in the natural language processing literature, my solution has been informed by my belief that the key to success in building systems that interact with human beings in natural language is to understand how humans process language. This dissertation is therefore heavily indebted to work in linguistics and psycholinguistics, and it is my hope that it can be seen as a contribution to these fields as well.

<sup>4</sup> These terms are all introduced below.

## 1.1 AMBIGUITY IN NATURAL LANGUAGE

### 1.1.1 Kinds of Ambiguity

*Lexical and structural ambiguities* are perhaps the best known kinds of ambiguities, and certainly the most studied by psycholinguists [Frazier and Fodor, 1978, Crain and Steedman, 1985, Small *et al.*, 1988, Grefen, 1989, Altmann, 1989, Marslen-Wilson, 1989]. One kind of lexical ambiguity occurs when elements of different syntactic categories have the same phonetic or written realization. For example, the string "can" can serve as the realization of both a modal auxiliary, as in (1.4a), and of a noun, as in (1.4b). A second kind of lexical ambiguity also exists, a lexical item of a certain category (e.g., a noun) may be associated with different senses. The noun "form" for example, has 24 different senses, and may denote, among other things, "a printed or typed document with blank spaces for insertion of required or requested information," as in (1.5a), or "a conduct regulated by extraneous controls (as of custom or etiquette)," as in (1.5b).

- (1.4) a I can speak English  
b You're opening a can of worms, my friend
- (1.5) a I'll give you a form to fill out.  
b It's not proper form for an ambassador to wear spandex tights at an official dinner

A sentence is *structurally* ambiguous when more than one *parse tree* or *s-structure* can be associated with it.<sup>5</sup> In the well-known (1.6a), for example, the prepositional phrase (PP) "with a telescope" can modify either the noun phrase (NP) "the hill" or the verb phrase (VP) "saw a man on the hill." Sentence (1.6b) (from [Hirst, 1987], p. 9) can be used either to answer the question "What are they doing?" (in which case "cooking" is interpreted as the head of a VP, or the question "What are these?" (in which case "cooking" is interpreted as an adjective modifying the noun "apples"). Structural ambiguities such as these may also result in a sentence having distinct semantic/pragmatic interpretations, in most semantic theories, this is a consequence of the fact that many *maximal projections* semantically denote *functions*, and the argument to which some of these functions (e.g., the functions that are the denotation of PPs) apply is specified by the parse tree's structure of the sentence, so that different s-structures correspond to different function/argument structures at the semantic level.<sup>6</sup> A sentence may of course be both lexically and structurally ambiguous, as shown by example (1.2a).

- (1.6) a I saw a man on the hill with a telescope  
b They're cooking apples

<sup>5</sup> These definitions are from Webster's 7th Collegiate Dictionary, Copyright © 1963 by Merriam-Webster.

<sup>6</sup> In the generative tradition (Chomsky, 1967, Chomsky, 1991, Haegeman, 1991), a sentence is characterized by a triple (d-structure, s-structure, LEPF). The d-structure encodes information about subcategorization, the s-structure represents the visible constituent structure, the Logical Form (LF) encodes information about the logical structure of the sentence such as quantifier scope, and the Phonetical Form (PF) encodes information about the sentence's phonetic realization.

<sup>7</sup> A maximal projection is the highest level of projection of a lexical head such as a noun or a verb. NPs, PPs and VPs are all examples of maximal projections [Haegeman, 1991].

<sup>8</sup> For an introduction to the idea that sentence constituents may denote functions see, for example [Towsky *et al.*, 1981, Chierchia and McConnell-Ginet, 1990].



Finally, scopal ambiguity results from the fact that the functional/argument structure of a sentence is not completely determined by its s-structure. Several phrasal elements denote functions whose arguments cannot be inferred from s-structure alone: they include quantified NPs such as "every student" or "most departments," modal auxiliaries such as "can" or "must," adverbs of quantification such as "always," and connectives such as negation, conjunction and disjunction. Following Heim [1982], I call the functions whose arguments cannot be determined by looking at the s-structure alone *operators*; examples of operators are functions such as quantifiers and negation.

An example of scopal ambiguity is shown in (1.7). The argument to the function denoted by the adverb of quantification "always" in (1.7a) may either be the object denoted by the sentence "The postman rings twice," resulting in the interpretation that in any situation the unique postman in that situation rings twice (shown in (1.7b)), or the object denoted by the sentence "x rings twice" where x denotes a certain postman (say, Bill Smith), as in (1.7c).

- (1.7) a The postman always rings twice  
b ALWAYS[THE POSTMAN(x)]x RINGS TWICE  
c THE x POSTMAN(x) ALWAYS [x RINGS TWICE]

Certain types of scopally ambiguous sentences, such as (1.8a) and (1.8b) are such that one of their interpretations is clearly preferred. In other cases, such as the sentences in (1.9), the preferences are weaker.

- (1.8) a We should hook up an engine to the buscar  
b She knows a solution to every problem  
c All that glitters isn't gold  
(1.9) a Every kid climbed a tree  
b Someone may bring a letter for me

Other types of ambiguity also exist. A sentence may be ambiguous in that a speaker may be performing different speech acts when using that sentence in different contexts; speech act ambiguity is exemplified by (1.2c). A sentence may be referentially ambiguous, since pronouns and other anaphoric expressions may in general be used in contexts that contain more than one antecedent. Other ambiguities, such as the so-called collective/distributive ambiguity, result from the use of plural quantifiers [Bunt 1985, Scha, 1981].

## 1.1.2 Ambiguity, Vagueness, and The Combinatorial Explosion Paradox

The definition of the term 'ambiguity' has been the object of much debate. Not everybody in the literature agrees that all the forms of 'ambiguity' I listed above are indeed cases of ambiguity; it has been argued that some of the cases of 'ambiguity' are really cases of 'vagueness'. Among these 'reductionist' proposals, of special interest for this dissertation is the one arguing that there isn't such a thing as 'scopal ambiguity' [Kempson and Cormack 1981]. I discuss this hypothesis below.

\*Examples of vague statements are such statements as "That car is red" or "Bill is tall".

Syntactic tests have been devised as a theory-independent means to classify a sentence as ambiguous or non-ambiguous; some of these tests are discussed by Zwicky and Sadock in [Zwicky and Sadock, 1975]. An example of these tests is the "do so" conjunction test developed by George Lakoff. Conjoining a 'truly ambiguous' sentence such as "They saw her duck" as in (1.10a), one does not get a 'crossed understanding' such that the referents of "they" saw Lee's pet water-fowl while George saw Lee knowing "and" himself. However, (1.10b)—whose first clause is 'ambiguous' between a reading in which John and Martha left separately, and one in which they left together—does allow for these crossed understandings. This suggests that the first sentence is indeed ambiguous, while the second isn't.<sup>10</sup>

- (1.10) a They saw her duck, and so did George  
b John and Martha left, and so did Dick and Pat.

The attempts to reduce ambiguity to vagueness are at least in part motivated by what I will call in this dissertation the *combinatorial explosion paradox*. Each form of ambiguity introduces the potential for a combinatorial explosion in the number of readings. Scopal ambiguity alone introduces a number of readings which grows with the factorial of the number of operators; thus (1.11), with 9 operators (not counting tense) should have at least 9! readings, without even considering lexical ambiguity (e.g., the ambiguity of "can") or structural ambiguity (e.g., the ambiguity as to the attachment of the PP "or almost every issue"). Yet, humans are usually able to interpret such sentences quickly and unconsciously. It is natural to wonder whether these sentences are really ambiguous. I'll argue in this dissertation that an explanation for this paradox can be found that does not require abandoning the conclusions on ambiguity reached in the field of linguistics.

- (1.11) A politicians can fool almost everybody on almost every issue most of the time, but he cannot fool everybody on every single issue all of the time (Hobbes)

Given that an agreed upon definitions of ambiguity and vagueness exist in the literature, I'll have to give my own definitions. The intention I try to capture can informally be described as follows: a sentence is *ambiguous* if its speaker may be meaning to make distinct statements about the world, depending on the context; whereas a sentence is *vague* if its speaker may only have one intention (i.e., there is only one statement about the world he may make using that sentence), but there exist more precise statements that refine the statement expressed by the speaker. For example, "They saw her duck" is ambiguous because its speaker may use it to express two entirely different statements about the world, whereas, if we assume that the lexical item "red" denotes a single concept, the sentence "That car is red" is vague under this definition because the statement expressed by that sentence it can be further specialized into it statements "That car is bright red" and "That car is magenta". These definitions can be made a bit more formal by using the notions of *propositional* or *situational type*<sup>11</sup> to characterize each statement that may be made by a sentence, as follows:<sup>12</sup>

<sup>10</sup>The boundaries of these tests are discussed by Zwicky and Sadock in the paper cited

<sup>11</sup>I use here both the terms 'proposition' and 'situation type' to indicate the characteristic function of a set of states of the world.

<sup>12</sup>Church and McCarthy-Gent propose that a sentence is semantically ambiguous if it is associated with more than one meaning in the language system. Van Eijk proposes in [Van Eijk, 1991] that the difference between

**Definition 1.1** A sentence is *semantically ambiguous* if it may be used by a speaker to express distinct propositions (situation types), and these propositions differ in the value they assign to at least one situation.

**Definition 1.2** A sentence is *semantically vague* if it expresses a proposition  $p$  that can be specified in distinct ways; that is, if there are at least two propositions,  $p_1$  and  $p_2$ , both of which entail  $p$ , but differ in the value assigned to at least one situation (i.e., one of them denotes a subset of the situations denoted by  $p$ ).

The definitions of semantic (sentence) ambiguity and vagueness just given are meant to be as simple and theory-neutral as possible; whether a sentence comes out as ambiguous or vague according to these definitions depends on the number of interpretations assigned to that sentence by a particular semantic theory.

It is important to note that according to my definition of ambiguity, an ambiguous sentence does not come out as equivalent to a disjunction of the distinct interpretations of the sentence. To treat an ambiguous sentence in such a way would be tantamount to propose that the sentence denotes a single proposition, namely, the proposition that is true at a situation if either of the distinct interpretations is true at that situation; according to my definition above, instead, an ambiguous sentence denotes a set of propositions. For example, according to the definition of ambiguity I gave above, the speaker of the sentence "They saw her duck" could either intend to say that the contextually established set of individuals denoted by the pronoun "they" saw a contextually specified female person lowering herself, or she could mean that "they" saw the pet waterfowl of that female person. According to the disjunction theory, instead, the speaker of that sentence had a single meaning in mind, albeit a disjunctive one; namely, she meant that it was either the case that "they" saw a contextually specified female person lowering herself, or it was the case that "they" saw the pet waterfowl of that female person. The two theories are only equivalent if something like the following axiom schema is assumed:

$$[A \text{ MEANS that } P] \vee [A \text{ MEANS that } Q] \equiv [A \text{ MEANS that } (P \vee Q)]$$

I'll add that the disjunction theory of ambiguity doesn't explain at all the combinatorial explosion paradox: if what humans do is to produce a big disjunction, that means they must at some stage consider all distinct interpretations.

Both sentences that are lexically ambiguous and sentences that are scopally ambiguous are classified as ambiguous under the above definitions. In fact, under some fairly reasonable assumptions about the relation between syntax and semantics, my definition of ambiguity subsumes the notion of structural ambiguity as well: a structurally ambiguous sentence will come out as semantically ambiguous, under my definition, as long as each distinct c-structure of the sentence corresponds to a distinct proposition, once the interpretation of the lexical items in the sentence has been fixed.<sup>13</sup>

ambiguity and vagueness is that the latter disappears when additional knowledge about the context of utterance is gained, while the former persists. Both definitions appear to be closely related to mine.

<sup>13</sup> This is a fairly reasonable assumption also according to Zwicky and Sadock (1975) and I know of no counterexamples. Note that it does not follow from the fact that two distinct c-structures correspond to distinct propositions – for example, the active and passive versions of a sentence may both denote the same propositions).

### 1.1.3 Ambiguity and Perceived Ambiguity

As it turns out, the distinction between semantic ambiguity and semantic vagueness plays, in the theory proposed in this dissertation, a less important role than another distinction, that between semantic ambiguity and perceived ambiguity.

**Definition 1.3** An utterance of participant  $B$  to participant  $A$  in a context  $C$  is perceived as ambiguous in  $C$  by participant  $A$  if  $A$  can find more than one interpretation for that utterance in context  $C$ .

In these situations,  $A$  will probably use the term 'ambiguity' to characterize its perception of  $B$ 's utterance, saying something like "You're being ambiguous." Considering that most natural language sentences are semantically ambiguous, this kind of statements ought to be extremely common—in fact, it should be uttered pretty much after every utterance. That this is not the case indicates that two distinct notions of ambiguity are needed, in the sense that we need to separate the traditional notion of semantic ambiguity from the notion of perceived ambiguity. Semantic ambiguity is a characteristic of most, if not all, sentences, but needs not be noticed by listeners (and in fact, often is only discovered by linguistic research); humans seem to be sensitive to perceived ambiguity, a property of utterances in context, that is only found in cases of miscommunication or whenever it serves a theoretical purpose, as in puns.

The distinction between semantic ambiguity and perceived ambiguity plays a central role in the theory developed in this dissertation, and in particular in the account of the combinatorial explosion paradox, the idea being that only perceived ambiguity results in a multiplicity of representations; semantic ambiguity is instead represented implicitly.

## 1.2 AVOIDING THE COMBINATORIAL EXPLOSION PARADOX

The combinatorial explosion paradox introduced in the previous section is seen by many as the central issue to be dealt with by theories of ambiguity and surface discourse interpretation. Researchers have been aware of this problem for quite some time, and a number of solutions have been proposed (Knech, 1975; Lapp, 1975; van Lehn, 1978; Wexler, 1978; Kempson and Cunnack, 1981; Scha, 1981; Fodor, 1982; Hobbs, 1983; Dunn, 1985; Fendall et al., 1987; Poesio, 1991; Alhawari, 1992; Kurtzman and MacDonald, 1993; Reyle, 1993). In this section I review this work, and I also introduce the approach that is currently believed by many (including myself) to hold the solution to the paradox.

### 1.2.1 Syntactic and Semantic Constraints

An important line of attack on the issue of combinatorial explosion is based on the observation that even though the number of permutations of the operators in a particular sentence may be rather large, constraints of a syntactic and/or semantic nature drastically reduce this number

- (1.15) a. A student from every department was at the party  
b. A student who was from every department was at the party

May proposed that the Scope Constraint holds because the scopally disambiguated interpretation of a sentence are represented at a (syntactic) level of representation called *Logical Form* (LF). The Logical Form of a sentence is obtained from its s-structure by means of a transformation called *Quantifier Raising*, and transformations are subject to a constraint called *Subadjacency*. This is the constraint that movement can cross at most one *bounding node* (the bounding nodes are S — R, in more recent terminology — and NP). The reading of (1.15b) in which "every department" takes wide scope over "a student" can only result from an instance of Quantifier Raising which violates Subadjacency.<sup>16</sup>

May [1983] also observed the asymmetry between subject and object positions in (1.19): (1.19a) is ambiguous, but (1.19b) isn't. These asymmetries are discussed in length by Chierchia in [Chierchia 1993].

- (1.19) a. What did every student buy for Max?  
b. Which student bought everything for Max?

Both Sag [1976] and Williams [1977] noted that sentences like "Some student admires every professor" are ambiguous in isolation, but in a VP-deletion context, they aren't.

- (1.20) Some student admires every professor, but John doesn't.

A lot of recent work—e.g., Kiss [1993], de Swart [1992] and Szabolcsi and Zwarts, 1997—concentrates on identifying similar constraints and exploring their nature. For example, Kiss observes that (1.21a) doesn't have a reading in which negation takes wide scope over the wh-phrase (contrast (1.21a) with (1.21b), which does have a 'family of questions' reading in which what is being asked if for each test, which student completed that test).

- (1.21) a. Which student didn't complete the test?  
b. Which student completed every test?

Any theory of scope disambiguation must clearly include an account of how these constraints operate and interact, which means that knowledge about syntactic structure and/or the semantics of lexical items must be brought to bear in the disambiguation process. It's also clear, however, that there is more to scope disambiguation than these constraints, because most of the examples

<sup>16</sup>The Scope Constraint is very strong. According to Van Lehn [van Lehn, 1978], who ran experiments in test which factors play a role on scope disambiguation (see below), 100% of his subjects agreed with the predictions of the Scope Constraint [1978]. Two classes of apparent counterexamples to the Scope Constraint are known, but the case is still open. Indefinite and definite descriptions do not seem to be subject to the constraint, as shown by (1.16). However, it has been argued that indefinite descriptions are not quantificational [Horn, 1982]. More trouble some cases like (1.17) involving quantifiers, a discussion of related examples is presented in [Peters and Zucchi 1992]. Finally, it should be noted that the predictions of the Scope Constraint are less robust when quantifiers contained in non-sentential complements are concerned, as in (1.18).

(1.16) Every student who wanted to believe found it charming.

(1.17) The woman that every Englishman respects the most is his mother.

(1.18) A quick test confirmed that each drug was psychoactive.

[van Lehn 1978]

Permutations may not correspond to actual readings for a least three reasons. First of all as noted by Hobbs and Shieber [1987], some permutations result in logical expressions that are either ill-formed or contradictory. For example, (1.12a), in the interpretation in which the pronoun "he" is anaphoric on the NP "every man," does not have a reading in which the NP "the woman he married" anticipates the NP "every man." There is no well-formed logical expression that may represent this reading.<sup>14</sup> Hobbs and Shieber point out that this constraint in general prevents a quantifier to scope between a noun and its complement. For example "a meeting" may not scope inside "most" but outside "each" in (1.12b).

- (1.12) a. Every man, loves the woman he married  
b. Most people on each committee attended a meeting.

Another reason why the number of actual readings of a sentence is much less than the number of permutations of its operators is that two distinct permutations may correspond to semantically equivalent readings. For example (1.13) has only one reading, even though (at least) two equivalent logical expressions can be obtained as the translation of the sentence.

- (1.13) A student saw a dog.

Finally, the readings corresponding to certain permutations may be unavailable because of semantic, or syntactic constraints. A very strong semantic constraint, extensively discussed by Labov [Labov, 1977], concerns the occurrence of polarity items such as 'any', 'or', 'yet', 'at least' [Labov, 1977], which is not enough for a polarity item to be in the semantic scope of a downward-entailing operator. The relevant examples are shown in (1.14): in the prepositional (1.14a) the polarity item "any buscar" is both in the semantic scope and in the command domain of the negation operator (1.14b), however, where the polarity item is only in the semantic scope of negation, is ungrammatical.

- (1.14) a. I don't believe that any buscar is at Avon  
b. \*I believe that any buscar is not at Avon

The work on uncovering missing readings has been especially intense in the generative tradition, where the emphasis has been on detecting readings that are absent due to constraints on syntactic transformations and/or conditions on syntactic levels of representation [May, 1977; May, 1985]. Some of the constraints proposed in this literature did not endure the test of time,<sup>15</sup> but others have been proved to yield quite robust predictions.

Among these latter, the best known example is perhaps the observation that a quantifier cannot take scope outside the clause in which it appears. The first mentions of "scope islands" can be found in Postal [Postal, 1974], Rodman [1976] and May [1977], the constraint was called Scope Constraint by Heim [1982]. The Scope Constraint is exemplified by the contrast in (1.15), whereas (1.15a) has a reading in which "every department" is allowed to take wide scope over "a student" this reading is not available for (1.15b), even though arguably "from every department" and "who was from every department" have the same derivation.

<sup>14</sup>Peters [1980] argues that this constraint is best formulated as a condition on semantic derivation rather than as a condition on the syntax of logical expressions.

<sup>15</sup>Two examples among many are the Linear Order constraint discussed below, according to which a quantifier always takes scope on all operators on its right, and what we may call the Past tense Constraint, according to which the subject NP in a passive clause always takes wide scope over other NPs in the same clause.

considered in this dissertation do allow for more than one interpretation. Consider, for example, the sentence *Afternoon* at the beginning of this chapter. "I can't find a piece of paper" clearly, we don't want a theory that rules out one of the two interpretations.

The proposal presented in this dissertation has been in part motivated by the need to account for the role that syntactic and semantic constraints play in the disambiguation process, but I haven't attempted either to uncover new constraints or to reinterpret those that are already known.<sup>17</sup>

## 1.2.2 Is There Such a Thing as 'Scopal Ambiguity'?

It has been noted in the literature that the syntactic tests for ambiguity do not classify a sentence as ambiguous if the propositions expressed by the sentence are such that one entails the other. This is the case, for instance, with sentences such as "Every kid climbed a tree" (Hakoff, 1970; Zwicky and Sadock, 1975; Kempson and Cormack, 1981).

This fact has been grounds for claiming that sentences such as (1.9a)—in fact, all quantified sentences—are not ambiguous. This would give us a simple explanation for the combinatorial explosion paradox: even though we would still be left with the task of accounting for the 'explosion' generated by other forms of ambiguity. The more articulated statement of this position has been made by Kempson and Cormack (K&C) in [Kempson and Cormack, 1981]. K&C propose that sentences like (1.23) are not ambiguous, but vague: according to them, such sentences semantically denote the weaker reading (the one in which the universal quantifier takes scope over the existential). The stronger reading is the result of pragmatic reasoning.

(1.23) Every linguistics student has read a book by Chomsky

What's more, Kempson and Cormack argue that this is also the case with sentences such as (1.24) that, because of the great number of different interpretations they may allow,<sup>18</sup> have traditionally been taken as problematic for ambiguity claims.

(1.24) Three examiners marked six scripts

In fact, K&C propose that all quantified sentences denote a single proposition, in this way the combinatorial explosion paradox disappears, at least as far as scopal ambiguity is concerned. This strategy of eliminating the paradox by claiming that quantified sentences have a unique,

<sup>17</sup> An important but not very well understood factor in disambiguation is played by sentence constituents sometimes called *scope disambiguation markers* such as "a different" or "the same": that when used to modify a determiner, are very effective in disambiguating the sentence.

- (1.25) a Every kid climbed a different tree  
b Every kid climbed a certain tree  
c Every horse entered the central  
d Every horse entered the same arena  
e I gave the kid a book, not a horse

<sup>18</sup> Hunt counted up to 30 readings for such sentences [Hunt, 1984].

'weakest' interpretation has been pursued by others as well, among them Hobbs [1981] and Kallman [1990].<sup>19</sup>

As we will see later, the solution to the paradox that I propose is also based on a processing strategy that does not require generating all of the disambiguated interpretations. Two aspects of K&C's proposal are, however, noticeable: these are the claims that there is no such a thing as 'scopal ambiguity', and that sentences with two or more operators denote the weakest proposition among the available translations.

The basic problem with K&C's proposal is that the observations on which it is based are only true of special cases of quantified sentences. In general, it is not true that a sentence with two quantifiers has two interpretations, one of which entails the other. (1.25) does not have an interpretation weak enough to be entailed by all others, yet able to capture the truth conditions correctly.<sup>20</sup>

(1.25) Few students know many languages

A second problem with the approach of Kempson and Cormack is that if one wants to claim that the meaning of a sentence such as (1.26a) is something like (1.26b), as Kempson and Cormack do, then one ends up predicting that the meaning of (1.26c) should be something like (1.26d), the strongest interpretation of the sentence. In other words, one either has to give up comprehensibility for sentences like (1.26b), or to abandon the strategy of letting sentences denote their weakest interpretation [Chierchia and McConnell Ginet, 1994].

- (1.26) a Every kid climbed a tree  
b  $(\forall x \text{ KID}(x) \supset (\exists y \text{ TREE}(y) \wedge \text{CLIMB}(x,y)))$   
c It is not the case that every kid climbed a tree  
d  $\neg \forall x \text{ KID}(x) \supset (\exists y \text{ TREE}(y) \wedge \text{CLIMB}(x,y))$

Finally, there is the question of preferences. I have noted above that most scopally ambiguous sentences have preferred readings in context, and I'll shortly discuss some psychological experiments that confirm this. Arguably, one could take Kempson and Cormack to predict that the weaker reading should always be preferred, there is evidence, however, that humans often prefer the stronger reading. For example, Kurtzman and MacDonald, whose experiments I discuss in detail below, observed that the preference for assigning to (1.27a) the reading in which 'a kid' takes wide scope (the stronger reading) is much stronger than the preference for assigning (1.27b) the weaker reading [Kurtzman and MacDonald, 1993]. Kempson and Cormack propose that interpretations are pragmatically filtered so that there is a preference for the interpretation

<sup>19</sup> The description of Hobbs' proposal, while broadly correct, doesn't make a complete justice. Space prevents a more detailed discussion.

<sup>20</sup> Actually, this problem arises already with the sentences involving numerical quantifiers that Kempson and Cormack looked at, such as (1.24). In order to derive the problem, K&C propose for these sentences a translation that is so weak as to fail to capture the correct truth conditions. More recently van der Dries and Verhey [Verhey and van der Dries, 1991] proposed a representation that they argue is entailed by all the different interpretations of (1.24) yet allows us to distinguish between different numerical quantifiers. In his dissertation, however, van der Dries argues that this representation is an improvement on the cumulative readings observed by Scha [Scha, 1981].

in which the 'topic' (that they define as the leftmost quantifier at logical form) corresponds with the syntactic subject. I'll return on this subject later.<sup>21</sup>

- (1 27) a A kid climbed every tree  
b Every kid climbed a tree

### 1.2.3 Underspecified Representations

I said above that while Kempton and Cormack's hypothesis that sentences with more than one operator are not semantically ambiguous is difficult to defend, their intuition about the direction that has to be taken to solve the combinatorial explosion paradox—humans have no troubles processing semantically ambiguous sentences because only a few disambiguated interpretations ever get generated—is rather plausible.

In fact, this intuition lies in well with the solutions to the problem of combinatorial explosion developed in the natural language processing (NLP) community. Schubert and Pelletier (1982), Hobbs (1982, 1983), Hobbs and Shieber (1987), the developers of TITAN (Grosz *et al.* 1987), Allen in his textbook (1987), Fenslad *et al.* (1987) and, more recently, the developers of the Core Language Engine (Althaus, 1992) all propose that the initial result of interpretation is a single (or a few)<sup>22</sup> underspecified representation—a representation that encodes some aspects of a sentence's meaning (typically, the result of syntactic interpretation and lexical disambiguation), but leaves other unspecified (e.g. the semantic scope of operators and/or the interpretation of anaphoric expressions). The underspecified representation serves as input to subsequent disambiguation steps that, in addition to resolving scope, are concerned with reference interpretation ellipsis, etc., and may make use of commonsense knowledge, general pragmatic principles, and so forth. For example, the representation for (1 28) proposed by Schubert and Pelletier is (1 29) in which quantifiers are left in place much as they would be at s-structure prior to Quantifier Raising.

- (1 28) Every kid climbed a tree  
(1 29) [<sub>every kid</sub> > climbed <a tree>]

The idea has been gaining consensus in recent NLP literature that the underspecification approach may in fact have cognitive plausibility, in the sense that the reason why combinatorial explosion is not a problem (or humans) is because humans are able to represent semantic ambiguity *implicitly* by means of an underspecified representation. Since this assumption plays an important role in what follows, I will give it a reverend name, *Underspecification Hypothesis*.<sup>23</sup>

**Underspecification Hypothesis (UH)** Humans are capable of representing semantic ambiguity implicitly by means of some form of underspecified representation.

<sup>21</sup> It would be possible to argue, perhaps, that the sentence reading is obtained by applying Kempton and Cormack's generalizing operations in the wider interpretation. Presumably this implies that Karttunen and MacDonald's subjects should take longer to process (1 27a) than to process (1 27b).

<sup>22</sup> Perhaps *only* a few such interpretation per parse tree—could be obtained in the case of structurally ambiguous sentences.

<sup>23</sup> As far as I know, this hypothesis is not explicitly advanced anywhere in the literature, although a number of authors seem to imply something like it.

The difference between the underspecification hypothesis I just presented and the approach proposed by Kempton and Cormack is as follows. According to Kempton and Cormack, (certain) sentences are not semantically ambiguous, and therefore only one interpretation exists. According to the UH, these sentences are ambiguous, but this ambiguity is not a problem because humans represent semantic ambiguity 'implicitly', so the semantic ambiguity of a sentence does not require its listener to produce a proliferation of interpretations. The underspecified representation approach, in other words, achieves the result desired by Kempton and Cormack—it does not require that all available alternatives be generated—but at the same time it allows us to retain linguistically motivated notions of ambiguity and vagueness.

The distinction between semantic ambiguity and perceived ambiguity discussed above enters into play at this point: according to the underspecification hypothesis, a sentence may well be semantically ambiguous, but an inference of that sentence is perceived as ambiguous only if the context suggests multiple interpretations to the listener; only in that case it may cause processing problems for humans.<sup>24</sup> (If we believe in the underspecification hypothesis, we need not worry if our semantic theory is 'readings-happy' (in use Verkuyl's (1992) felicitous expression), as long as our discourse interpretation theory does not require us to generate them all.)

This of course means that the UH by itself does not suffice to solve the Combinatorial Explosion Paradox; a theory of discourse interpretation that does not assume that all interpretations are actually generated<sup>25</sup> is also required. In fact, a large part of the dissertation is dedicated to developing a theory of scope interpretation that, while assigning to an ambiguous sentence a set of meanings, does not require 'a whole set to be generated explicitly, and therefore doesn't suffer from the problems attributed by M&C to what they call the 'maximal ambiguity theories'.

The underspecified representations used in most NLP systems are treated as uninterpreted data structures: scope disambiguation in these proposals is therefore handled by means of extra-logical procedures such as those developed by Hobbs and Shieber (Hobbs and Shieber, 1987) or Harum (Harum, 1988). It is also known, however, that the choice of an interpretation often depends on the result of reference interpretation processes and/or on inferences based on commonsense knowledge, and therefore these processes and inferences should be integrated with the processes dedicated to scope disambiguation. For example, 'the cat' is assigned narrow scope in the preferred interpretation of (1 34) because that NP is interpreted as anaphoric on the NP 'a cat' in the antecedent of the conditional. In (1 35), world knowledge about ownership relations and the typical size of backpicks force the interpretation in which 'every student' takes wide scope. Anaphoric and commonsense knowledge (about affairs and about working relations) conjoin in making preferred the interpretation of (1 36) which requires more than one secretary. Knowledge about recent events in Italy makes the reading of (1 37) in which a single scandal is discussed most salient. Finally, the preferred reading of (1 39a)<sup>26</sup> is the one

<sup>24</sup> The very examples of 'clearly non-ambiguous sentences' discussed by Zwicky and Sadock are a case in point. All of these sentences may be perceived as ambiguous in the appropriate context: (1 30) becomes ambiguous in a context in which the speaker has two sisters and the hearer is aware of that, (1 31) becomes ambiguous if the order of the sisters is relevant—for example, if the statement is made to a police officer while identifying a criminal (1 30) My sister is the Raritanian head of state.  
(1 31) She wore a vest.

<sup>25</sup> This nice maximal pair is due to Kurt Van Lehn.

where "a LISP dialect" takes wide scope, whereas the preferred reading of (139b) is the one where 'everyone at ILLAC' takes wide scope; the contrast clearly depends on world knowledge < factors alone (namely, the fact that Xerox PARC researchers developed a LISP dialect called INTERLISP)<sup>26</sup>

- (134) If a cat sees a . . .  
 (135) Every student carries his books in a backpack.  
 (136) Every conservative MP has had an affair with his secretary.  
 (137) Most Italian politicians are involved in a scandal of enormous proportions.  
 (138) Every conservative MP has an office.  
 (139) a. Everyone at PARC knows a dialect of LISP.  
 b. Everyone at ILLAC knows a dialect of LISP.

In some of the literature on underspecified representations starting from (Schubert and Pelletier, 1982) and (Fenshau *et al.*, 1987), one can find, in addition to a clearer formulation of the underspecification hypothesis, attempts to spell out the consequences for 'model-theoretic' semantics of this hypothesis, i.e. to define more precisely what an ambiguous sentence denotes. A further step is taken in work such as (Pierro, 1991; van Lier, 1991; Alami and Crouh, 1992; Reyle, 1993) where logics are developed that include underspecified representations either as special kinds of terms or as special kinds of 'worlds', and thus can be used to describe such (communicative) contexts and the processes involved in scope assignment and reference disambiguation.

At the moment, the only proposals that appear capable of dealing with the 'Combinatorial Explosion' Paradigm are based on the version or the other of the Underspecification Hypothesis. I adopt it as well. In this dissertation, I simplify matters by assuming that both structural and lexical disambiguation take place before scope disambiguation and reference interpretation, so that I will not have to worry about encoding structural and lexical ambiguity in the underspecified representation I use; this representation will be an  $\lambda$ -structure whose leaves have been replaced by unambiguous semantic translations of lexical items.<sup>27</sup> The model-theoretic interpretation of

<sup>26</sup> The influence of world knowledge on the interpretation of ambiguous expressions is also easy to see. In (143) for example, the pronoun 'it' is clearly interpreted as referring to 'the audio part' even though two references are available, and the choice seems to depend on the fact that it is part of the back of computers. In the (rightly edited) fragment of an actual conversation in (132), 'the empty house' is interpreted as referring to the house mentioned in (131a), but this clearly requires that B makes the inference that A has not yet seen empty after what is said in (131a).

- (142) Hand up the cable to the audio part. It's on the back of the server.  
 (143) a. A: move the base to Aom and under the oranges.  
 b. B: ok.  
 c. A: then, and the empty house to Matt.

<sup>27</sup> Of course, there is plenty of evidence that at least some lexical disambiguation and reference interpretation operate in parallel (see, e.g., (132)). (Alami and Neeleman, 1988) is doing an empirical study on structural and referential disambiguation in carrying out different steps. The only problem is to develop a more complex form of underspecified representation that would allow for this.

ambiguity I will postpone to the definition of semantical ambiguity discussed above, according to which a sentence is semantically ambiguous if it is associated with a set of meanings of cardinality greater than 1.

## 1.2.4 Beyond the Underspecification Hypothesis

One might wonder whether the Underspecification Hypothesis perhaps in conjunction with the existence of syntactic and semantic constraints, explains the Combinatorial Explosion Paradigm. After all, if the process of interpretation stops after an underspecified representation is obtained, no explosion of the number of readings will result. We must ask ourselves the following questions: does further processing take place in humans after an underspecified representation has been obtained? And even if that is what occurs in humans, should the same approach be followed when developing a natural language processing system? I propose that the answer to both questions should be a qualified 'yes', and therefore the UH by itself is not sufficient to explain the CEP.

I will call the hypothesis that no disambiguation takes place once the underspecified interpretation has been obtained, 'un-': perhaps there are some reasons for doing so (e.g., 'new information is obtained', 'lazy disambiguation proposal'). There are several kinds of evidence suggesting that the lazy disambiguation proposal is not a plausible account of discourse interpretation, i.e., that humans do not take a lazy approach to discourse interpretation. First of all, there are psycholinguistic experiments indicating that whenever human subjects are presented with a sentence as part of a task, they arrive at a disambiguated interpretation, whether the sentence being processed is lexically, structurally, referentially or scopally ambiguous. In fact, in the case of the first three kinds of ambiguity at least, it's commonly accepted that disambiguation not only occurs, but it takes place rather early, as shown most clearly by phenomena like garden-path sentences (Frazier and Fodor, 1978; Craun and Steedman, 1985; Small *et al.*, 1988; Gordon, 1989; Altmann, 1989; Marsten-Wilson, 1989)<sup>28</sup>. In the case of scopal ambiguity, we will see later, there is evidence that preferences exist, it's not clear when disambiguation takes place (Kurtzman and MacDonald, 1993), but effects similar to garden path can be had with quantifier scope as well, as shown by (140) (Barwise [1987] calls these examples *jungle paths*)<sup>29</sup>.

- (140) Statisticians show that every 11 seconds a man is mugged here in New York City. We are here today to interview him.

Now should we forget the fact discussed above that humans are clearly able to recognize ambiguity in context when it occurs. This is shown both by the fact that the sentence's perceived ambiguity can be exploited for rhetorical effects, and by the fact that when clarity is a goal, people tend to construct the sentences occurring in natural conversations and texts in such a way as to

<sup>28</sup> Ambiguous expressions are the most clear indication that interpretation does not stop at the level of underspecified representations when reading texts like 'John met Bill. He was wearing a spinning jacket' most readers choose a decision about the antecedent of the pronoun 'he' or find the text ambiguous.

<sup>29</sup> Thanks to David McDavid for pointing out these examples to me.

avoid ambiguity so that most sentences one runs across in real texts or transcripts of natural conversations have preferred readings.<sup>30</sup>

All of this suggests that in general underspecified representations are not the ultimate result of discourse interpretation: further processing takes place, and this processing results in a single interpretation, in the case of (contextual) non-ambiguity, or, in more than one, if an ambiguity is perceived (whether or not that ambiguity was intentional). The only question that makes sense to ask is *when* disambiguation occurs, even though all evidence indicates that the participants in task-oriented conversations attempt to disambiguate their utterances right away, this question ultimately can only be answered by psycholinguistic research.<sup>31</sup> Whatever the answer to this question is, an account of disambiguation is needed even if our only goal is to solve the Combinatorial Explosion Paradox.<sup>32</sup>

If, in general, a speaker intends his listener to do more than simply arrive at an underspecified interpretation, a theory of ambiguity must include an account of the processes by which the intended interpretation is obtained, how these processes may result in a perceived ambiguity, and why combinatorial explosion is not a problem. That's because the Underspecification Hypothesis is also consistent with a theory of discourse interpretation requiring, say, that once an underspecified interpretation is obtained, all possible disambiguated interpretations have to be generated as a matter of fact: most theories of discourse interpretation developed in the natural language processing literature are of this kind. The UH thus needs to be supplemented with a theory of disambiguation that makes the process that results in the final interpretation(s) highly constrained.

In other words, we need, in addition to the Underspecification Hypothesis, a further assumption, that I will call *Anti-Random Hypothesis*.<sup>33</sup>

**Anti-Random Hypothesis (ARH)** Humans do not randomly generate alternative interpretations.

<sup>30</sup> Most people also tend to find the examples of scope ambiguity usually presented in the literature – in which the author (usually) provided all possible, but, rather artificial. This point has been impressed on me most forcefully by (George Hirst (p. 1) and Len Schubert (p. 1)).

<sup>31</sup> I am allowing for the possibility that the processes dedicated to scope assignment are of a different nature than those concerned with parsing and lexical disambiguation, both of which appear to be occurring on-line, as shown, for example, by the phenomenon of garden-path sentences. As I will discuss below, the procedure for assigning a scope to operators described below abstracts away from the characteristics of these other interpretive processes.

<sup>32</sup> This does not exclude the possibility that the disambiguation may be incomplete when the task at hand does not require complete understanding (as in rote memorization) or requires that the ambiguity be preserved (as in translation). Van Diermen argues rather convincingly that no sentence ever gets completely disambiguated (van Diermen 1991). It has also been argued that, for applications such as information retrieval or machine translation, a complete disambiguation might be useless or even incorrect (say, in case the original text is intentionally ambiguous) (Alchourrón 1992, Hirst 1991). We want to allow, therefore, for the final interpretation to be underspecified at least in certain circumstances. This is an additional motivation for assuming that the information provided by underspecified representations may appear in inference patterns. Even if we consider translation, however, some steps of interpretation appear necessary. Consider for example the task of translating English into Italian, a language in which adjectives agree in gender with the noun they qualify. To translate in Italian the following strip of a script:

(141) Cut the eggplant in thin slices with a knife. Chop a very sharp one  
to translate "sharp" into Italian it is necessary to understand its referent, since the translation of "eggplant" in Italian is *melanzana* (feminine) while the translation of "knife" is *coltello* (masculine).

<sup>33</sup> Although this assumption has never, to my knowledge, been made explicit in the literature, it is implicit in proposals such as Schubert and Huang's (Huang and Schubert, 1993) or Hobbs' work on abduction.

ons of an ambiguous sentence: only those few interpretations are obtained that (i) are consistent with syntactic and semantic constraints and (ii) are suggested by the context, except perhaps when no interpretation may be obtained by 'normal' means, in which case some sort of 'puzzle mode' may be entered.<sup>34</sup>

### 1.3 TOWARDS A THEORY OF SCOPE DISAMBIGUATION: DISAMBIGUATION FACTORS

Having clarified the concepts of semantic ambiguity, perceived ambiguity and of vagueness, and having argued that two hypotheses need to be made in order to account for the Combinatorial Explosion Paradox, I can now turn to the main task of this dissertation, to produce a theory of scope disambiguation consistent with these definitions and these hypotheses.

I review in this section the existing literature on scopal ambiguity processing. I will observe that a lot of principles have been proposed to account for the data about scoping preferences (Lakoff, 1971; Joup, 1975; van Lehn, 1978; Fodor, 1982; Kurtzman and MacDonald, 1993), but there isn't a theory explaining why do these principles play a role, how they operate, and how conflicts between principles are resolved.

#### 1.3.1 Linear Order

Most proposals on scope disambiguation<sup>35</sup> were developed to account for the general preference of the leftmost quantified phrase to take wide scope in simple active sentences like (1.42).

(1.42) Every kid climbed a tree. [Kurtzman and MacDonald, 1993]

G Lakoff (1971) proposed that this preference is due to the fact that sentences are parsed from left to right, "every kid" takes scope over "a tree" because it is processed first. (Kurtzman and MacDonald (1993) call this the *linear order principle*.) Lakoff actually argued that this principle effectively disambiguates (1.42), which, according to him, has only the reading in which "every kid" takes scope over "a tree." Fodor's proposal (Fodor, 1982) is also based "on the working assumption that these preferences are due to sentence comprehension processes" (p. 143).

#### 1.3.2 Grammatical Function, Quantifier Hierarchy, and Topic

In the first part of her paper (1975), Joup argues that "... in natural language, order has little to do with the determination of quantifier scope." (1975, p. 37). Joup presents examples from English and 13 other languages in which the scope assignment is opposite to that predicted by

<sup>34</sup> This entails that surface discourse interpretation may be different from the process of lexical disambiguation, where there is evidence that all interpretations are considered prior to context filtering (Tanenhaus et al., 1979; Hirst, 1987; Geurts 1989; Scott et al., 1988).

<sup>35</sup> More explicitly Fodor's (Fodor 1982).

the linear order principle. The preferred reading of (1 43), for example, is the one in which the NP 'each child' takes wide scope.

- (1 43) I saw a picture of each child [Loup, 1975]

The second part of the paper is devoted to the presentation of Loup's own proposal. According to Loup the relative scope of quantifiers is determined by the interaction of two factors. First of all, quantifiers whose determiner is "each" or "the" have the inherent (i.e. lexical) property of taking wide scope over indefinites, which in turn are lexically marked to take scope over plural quantifiers like "all". This hypothesis is motivated by contrasts such as those in (1 44)-(1 45), and also accounts for cases such as (1 43).<sup>36</sup>

- (1 44) a I saw a picture of each child  
b I saw a picture of all the children  
(1 45) a Fihel has a dress for every occasion  
b Fihel has a dress for all occasions

Loup proposes that a second hierarchy exists, among grammatical functions, such that listeners tend to attribute to NPs in subject position wide scope over NPs in indirect object position, which in turn tend to take wide scope over NPs in object position. The hierarchy between grammatical functions accounts for the preferred reading of (1 42), and apparently holds for all languages Loup analysed. In Greek, for example, the direct and indirect object c in the mismatched reversing the order of the quantifiers without however affecting the preferred meaning. So the preferred reading of both (1 46a) and (1 46b) is the one in which a unique cake is being looked for all the girls irrespective of the linear order of the NPs, "na glyko" and "pita kapota korisia". Loup discusses similar examples in Turkish and Tagalog.

- (1 46) a Giannis ekane ena glyko gi kapta korisia  
John made a cake for some girls  
Taro made a cake for some girls  
b Giannis ekane gi kapta korisia ena glyko  
John made for some girls a cake  
John made a cake for some girls

Loup also observes that the NP expressing a sentence's topic tends to take wide scope. Topic is not explicitly marked in English, but in languages in which topic is marked, like Japanese or Korean, the preference is rather clear. (1 47b) is ambiguous, but the reading in which the NP in subject position, "most students" takes scope over the NP in object position, "every language", is preferred. This preference is maintained if the NP in object position is scrambled in sentence-initial position, as in (1 47c) (another counterexample to Lakoff's linear order principle). If however, the NP is marked with the topic-marking suffix "wa" as in (1 47d), so clearly the preferred reading of the sentence becomes the one in which "every language" takes wide scope.<sup>37</sup>

<sup>36</sup> Hunka [1974], Krach [1974], van Lén [1978], Hendrix [1978] and Menn [1988] also study the effect of lexical preferences or weights as they are also called.

<sup>37</sup> Thanks to Leonidas Krachinades for confirming the judgments reported by Loup and correcting the sentences.

<sup>38</sup> The examples are mine; the judgments were verified with the help of five native speakers.

- (1 47) a Most students speak every language.  
b Hotoondo no gakusei-ga subete-no gengo-o hanasu  
most-gen student-nom every language-acc speak  
c Subete-no gengo-o hotoondo-no gakusei-ga hanasu  
every language-acc most-gen student-nom speak  
d Subete-no gengo-wa hotoondo-no gakusei-ga hanasu  
every language-TOP most-gen student-nom speak

A weak topicity effect may also be had in English, where the NP in subject position often plays the role of sentence topic, and special constructs such as left dislocation and fronting have the purpose of indicating a sentence's topic [Gundel, 1974; Reinhart, 1981].

- (1 48) a As for this room, it really depresses me  
b In this room, I feel really depressed

Loup proposes to account for the data about topic by introducing a new grammatical function for Topic, and by revising the hierarchy of grammatical functions so that Topic takes precedence over Subject.<sup>38</sup>

### 1.3.3 Scope and C-Command

Jackendoff [1972] and Reinhart [1976, 1983], ch. 3 and 9) propose to account for the preferred reading of (1 42) by stipulating that a quantified expression takes scope over another quantified expression if the latter is in the c-command domain of the former at s structure in (As in the case of Lakoff, Jackendoff and Reinhart actually claim that these sentences are unambiguous). This generalization accounts not only for the preferred reading of (1 42) but also for the preference for NPs in fronted pps to take wide scope, exemplified by the following contrast (the data are from [Reinhart, 1983], p. 192).

- (1 49) a Some reporters worship Kissinger in every town he visits  
b In every town he visits, some reporters worship Kissinger

In (1 49a), there is a slight preference for "some reporter" to take wide scope, but the other interpretation is available. In (1 49b) there is a clear preference for "every town" to take wide scope.

<sup>39</sup> Kar [1980], p. 26) argues that the process that determines the relative scope of operators relies on pragmatic information only. He proposes that by adopting a Topic Principle according to which the NP in topic tends to take wide scope we would account not only for the facts about topic observed by Loup, but also for the preferred reading attributed in sentences like (1 42), the idea being that the preference for Subjects to take wide scope would be a consequence of the fact that Subjects in English are normally used to express the topic of a sentence. It is not clear how this principle would account for the other facts explained in terms of the grammatical function hierarchy. The topic position was independently proposed by Kampen and Crumack [1981].

<sup>40</sup> Karttunen and MacDonald [1991], all the the C-command principle.



### 1.3.4 Van Lehn's Results

in his MS thesis (van Lehn, 1978), van Lehn's looked at the impact on scope disambiguation of constraints on transformations such as the Scope Constraint, as well as of the principles proposed by Inup (Quantifier Hierarchy and Grammatical Function Hierarchy).

Van Lehn selected 121 sentences out of technical papers written at MIT. Each sentence was typed on a file card, and submitted to the informant to be read silently to avoid the effect of intonation. The informants were first asked to paraphrase the sentence, if that alone didn't indicate a clear preference, questions were asked of the informant. For example, given the sentence "Every guy kissed a girl", questions such as "Did they all kiss the same girl?" were asked.

Van Lehn did find evidence for a Quantifier Hierarchy, albeit not of the form proposed by Inup. More precisely, he found that in sentences with an indefinite NP and a universal quantifier or a plural quantifier, 100% of the informants preferred the reading in which the indefinite NP took narrow scope when the determiner "each" was used; the preference became lower and lower as other quantifiers were used and, with all NPs, the preference was reversed. In (1.50a), there is a clear preference for "a glass of champagne" to take narrow scope. In (1.50b), there is no clear preference, and in (1.50c), there is a strong preference for the NP "a glass of champagne" to take wide scope.

- (1.50) a The club president splashed each member with a glass of champagne  
b The club president splashed many members with a glass of champagne  
c The club president splashed all members with a glass of champagne

Van Lehn noted that, considering only NPs with universal force, the likelihood that a universally quantified NP would take wide scope was inversely correlated with the acceptability of that NP as the subject of collective predicates such as "met".

- (1.51) a \*Every man met.  
b \*Every man met.  
c ??All of the men met.  
d All the men met.  
e The men met.

Inup had suggested in her paper that a direct correlation exists between the position of a quantifier 'D' P' (where D is the determiner, P a property) in the quantifier hierarchy and the ratio between all P's and the percentage of P's that have to have the property Q in order for the statement "D P's are Q's" to hold. For example, every quantifier with universal force would be higher in the quantifier hierarchy than a quantifier such as 'most P's'. Van Lehn's data suggest that the position in the quantifier hierarchy has nothing to do with the relative size of the set. He proposed to replace Inup's version of the quantifier hierarchy with a *distributivity hierarchy* between quantifiers with universal force: the higher a determiner is in the hierarchy, the likelier it is that that determiner is going to take wide scope. Van Lehn's distributivity hierarchy is as follows:

each < every < all < the (plural)

Furthermore, on the basis of examples like (1.52), indicating that a clear contrast exists between the scoping behavior of definite and indefinite NPs, van Lehn proposed that a distinct principle was involved, according to which specific NPs (he considers definite NPs specific) take wider scope than non-specific NPs.

- (1.52) a The club president splashed each member with a glass of champagne.  
b The club president splashed each member with the glass of champagne.

Van Lehn also studied the relation between clause type and the availability of wide scope reading, both in relative clauses and in sentential complements. He did note a strong effect, although not quite as strong as the effect of the distributivity hierarchy and the specific/non-specific distinction. For example, he noted the contrast in (1.53) in (1.53a) the reading in which "Each raw rubber producer in Brazil" takes wide scope is not available, while in (1.53b) that reading is as likely as the reading in which "a guy" takes wide scope, and in (1.53c) that reading is preferred by 100% of Van Lehn's informants.

- (1.53) a At the conference yesterday, I managed to talk to a guy who is representing each raw rubber producer in Brazil.  
b At the conference yesterday, I managed to talk to a guy representing each raw rubber producer in Brazil.  
c At the conference yesterday, I managed to talk to a representative from each raw rubber producer in Brazil.

Van Lehn proposes an *embedding hierarchy* clause < gerund < PP (in account for these data the higher a clause is in the hierarchy, the more likely it is that a quantifier in it can take wide scope over another quantifier in the same sentence).

Finally, van Lehn tried to verify whether the position of an NP at s-structure affects the scope that that NP is going to take. He was especially interested in comparing the predictions resulting from the Linear Order principle with the predictions based on Inup's Grammatical Function Hierarchy, as he noted that these predictions only differ in a couple of cases (e.g., the linear order principle predicts that NPs in object position take scope over NPs in adverbial PPs). His results were inconclusive: both principles appear to be in good agreement with most of the preferences, and in disagreement with the preferences observed in about a third of the cases. He suggests that, this being the case, surface order may be preferable because theory-independent and easier to implement.<sup>41</sup>

### 1.3.5 The Use of Disambiguation Factors in NLP Systems

Missing from the literature reviewed so far is an account of how the proposed disambiguation principles may result in an interpretation, and how they interact. This has been a central concern

<sup>41</sup> The hierarchy between clause-types that would be obtained by assuming the s-command principle also differs only in few respects from the hierarchy resulting from linear order. The 's-command hierarchy' would be as follows:  
subject >  
prepositional pp and topicalized pp >  
adverbial pp and adverbial np > verb phrase pp > object

for implementors of NLP systems, who have proposed several architectures, all based on the idea of choosing a number<sup>47</sup> of disambiguation factors, assign them a relative weight, and use them to score the readings that are compatible with a set of syntactic and semantic constraints.

Let's consider the procedure adopted in TEAM (Grosz *et al.*, 1987), for example. The constraints the system knows about include the Scope Constraint and a constraint preventing the generation of logical forms that contain unbound variables. The disambiguation factors chosen by the developers of TEAM include the Linear Order principle and the Quantifier Hierarchy. The procedure for generating the preferred interpretation consists of the following steps:

- 1 Generate all the readings consistent with the constraints.
- 2 Assign each reading a score depending on the number and the weight of the disambiguation factors it satisfies (say, 1 point if it satisfies the quantifier hierarchy principle, multiplied by the weight of that principle, 1 point if it satisfies the linear order principle, and so forth).
- 3 Pick the reading with the highest score.

This strategy is potentially subject to combinatorial problems, but in practice only a few readings at a time would ever be generated by TEAM, due to the nature of the interaction with the system (a translatable interface to data bases). The real problem for this approach to scope disambiguation derives from the lack of any theoretical understanding of the nature of the scope disambiguation process: without such an understanding, the designers have no way of choosing among the many existing principles, or of assigning them a relative weight, except then by trial and error. After the designer of a system has spent a long time choosing the principles and fine tuning their relative weights so as to get the correct predictions for the training corpus, there is no guarantee that the principles or the weights used for one application will give good results when different data, a different domain, or a different kind of interaction are considered.<sup>48</sup>

### 1.3.6 Kurtzman and MacDonald's Experiments

Luckily, some insights into the workings of the scope assignment process can be obtained by the one study in the recent psycholinguistic literature concerned with scope disambiguation. This study was conducted by Kurtzman and MacDonald (henceforth K&MD), who, like van Lehn, set out to test the predictions of the scope disambiguation principles proposed in the literature, but relied on the experimental techniques developed in the psycholinguistic literature to get a better control of the experimental setting (Kurtzman and MacDonald, 1993).

Kurtzman and MacDonald designed their experiments so as to test the predictions of the Linear Order Principle, the Command Principle, one of the two principles proposed by Inup (Grammatical Function Hierarchy), and of two principles not explicitly proposed in the literature

<sup>47</sup> Usually many more than three per discussion, see e.g. [Meeus 1988].

<sup>48</sup> The abstraction method developed by Hobbs and colleagues [1990] and the Rayman-Nick technique for language interpretation proposed by Charniak and Goldman [1988], although more elegant than the procedure just described, are based on the same principle and hence are subject to the same criticism.

but consistent with observations made elsewhere, namely, the topic principle according to which topics tend to take wide scope, and the thematic hierarchy principle according to which NPs that fill roles higher in the thematic hierarchy (Jackendoff, 1972; Grimshaw, 1990) are more likely to take wide scope.<sup>49</sup>

In K&MD's experiments, subjects sitting in front of a computer terminal, after pressing the space bar, were presented with a first sentence, the *quantified phrase*, that contained two quantifiers. This sentence would then disappear when the subjects pressed the bar again. To be replaced by a second sentence, the *continuation*, chosen in such a way as to be easier to interpret if one of the interpretations was assumed (an example of quantified phrase and two continuations is shown in (1.54)). The subjects had to decide whether the continuation was plausible, given the context provided by the quantified phrase. An interpretation was considered preferred when the percentage of subjects finding the continuation consistent with that interpretation plausible was larger than 50% and significantly larger than the percentage of subjects that found the other continuation plausible.

- (1.54) Every kid climbed a tree.  
(1.55) a The trees were full of apples  
b The tree was full of apples

Kurtzman and MacDonald ran four experiments, each of which involved around 40 subjects. Each subject was presented with both ambiguous and unambiguous quantified sentences, e.g., "Every kid climbed the same tree" or "Every kid climbed a different tree" would be used in addition to "Every kid climbed a tree". Other variables tested in all experiments were the influence of the quantifier order (e.g., "A kid climbed every tree" was used in addition to "Every kid climbed a tree") and of the choice of predicate, both 'action' and 'perception' predicates such as "see" were used. In the first experiment, active sentences were used, in the second passive sentences, in the third and fourth experiments, one of the quantifiers was embedded in a clause modifying the other quantifier, as in "A picture of every admiral".

K&MD found, first of all, that in a number of cases a continuation was clearly preferred over the other. For example, they found that indeed there was a preference for the subject in active sentences to take wide scope. Interestingly, the preference was stronger when the quantified phrase was of the form "A P" (as in "every") than when it was "A boy climbed every tree",<sup>50</sup> than when it had the form "Every P X" (as in "every") and in "A picture of every admiral" would predict.<sup>51</sup>

In fact, the second important finding of Kurtzman and MacDonald was that all of the principles they looked at were subject to exceptions, for every principle proposed in the literature they found cases in which the preferences of their subjects were in disagreement with the predictions of the principle. This is important because claims have been made in the literature

<sup>49</sup> The notion of thematic role, as well as the thematic hierarchy, are discussed in more detail in Chapter 6.

<sup>50</sup> In this case 85% of the subjects found plausible the continuation consistent with the reading in which the subject NP took wide scope, while only 25% of the subjects found plausible the one in which object NP took wide scope.

<sup>51</sup> The percentage in this case were 75% and 15% respectively.

<sup>52</sup> Kurtzman and MacDonald tentatively propose a single preference principle to account for these data, according to which an adverbial or subject or topic position is immediately interpreted. See below.

that the *sc* principles made sentences unambiguous. For example, K&MD observed a preference for the embedded NP in a complex nominal take wide scope, the preferred reading of (1.56) would thus be the one in which 'an admiral takes wide scope over 'each picture'. This result is not predicted by any of the principles tested by Kurtzman and MacDonald.<sup>40</sup>

(1.56) George owns each picture of an admiral

Agreement appeared to affect the preference for subjects in active sentences to take wide scope. The preference was stronger when the subject of the verb had the characteristics of an agent as in (1.57a), than when it had the characteristics of an experiencer, as in (1.57b).<sup>41</sup> This result would indicate that a principle formulated on the basis of thematic roles, such as the Thematic Hierarchy principle, may be account for the preference for 'every kid' to take wide scope in (1.57a).

- (1.57) a Every kid climbed a tree  
b Every kid saw a tree

Finally, K&MD observed a clear contrast between active and passive sentences. While their subjects showed a preference for the reading of active sentences such as (1.58a) in which the subject NP takes wide scope, no such preference was observed for the passive version of the same sentence.<sup>42</sup>

- (1.58) a Every winter wrote a book  
b A book was written by every winter

Most interesting is Kurtzman and MacDonald's account of this contrast: they do not propose a new 'principle' to account for these data, but argue that what is being observed is a conflict between two principles, a principle assigning wide scope to subjects and the principle assigning wide scope to agents. Although Kurtzman and MacDonald do not mention it, further support for this proposal is provided by well-known data about passives found in the literature, sentences such as (1.59). The main predicate in these sentences is stative, which means that no conflict of the type observed by Kurtzman and MacDonald in sentences such as (1.58b) should be observed in fact: the reading where the subject NP takes wide scope is much preferred.

- (1.59) Two languages are spoken by everyone on Cornman Islands

Kurtzman and MacDonald caution against reading too much into their results. They argue that the factors that determine the choice of the preferred reading cannot be identified on the basis of their data alone (1993, p. 40). One should also keep in mind that only sentences with simple quantifiers such as 'every' and 'a' were considered. Nevertheless, their results lead to a perspective on the factors affecting scope disambiguation that is rather different from those

<sup>40</sup> Examples like (1.56) were considered by May in his thesis (1977) and by Reinhart (1997).

<sup>41</sup> While in the first case the preference for the two readings in which the universal NP and the indefinite NP take wide scope were 71% and 33%, in the second case the preferences were 75% and 67% respectively.

<sup>42</sup> The percentages of subjects favoring the first and second reading (passive) were 57% and 45% respectively. Interestingly, on the same percentages were observed when an indefinite NP occurred in subject position. This result, even more double on the simple reference principle proposed by Kurtzman and MacDonald in account for the disputes observed in active sentences.

proposed in the previous literature. According to this new view, the 'principles' are always active, but they behave as *pragmatic defaults* that may be overridden by stronger defaults or originate conflicts with defaults with the same strength. Kurtzman and MacDonald also hypothesize that '... processes that are not strictly dedicated to the interpretation of scope relations may nevertheless influence the interpretation of quantifier scope ambiguities.' (1993, p. 22). This observation is motivated by their results about the significant difference in preference for the subject NP to take wide scope between the case in which the subject NP is indefinite and the case in which it is an every NP. The theory of scope disambiguation and discourse interpretation I propose is based on these ideas.

As far as the disambiguation factors themselves are concerned, three main facts emerge from K&MD's experiments: the preference for subject NPs to take wide scope in active sentences, the rule played by agentivity, and the little evidence for a Linear Order principle.<sup>43</sup>

Kurtzman and MacDonald conclude that "... the results leave open the question of whether the building and selection of representations of scope are mandatory processes" (1993, p. 45).

### 1.3.7 Summary

Let me summarize briefly the material presented in this section. Certain factors about scope disambiguation have been repeatedly observed looking at sentences in isolation. Both formal and informal experiments have been conducted these experiments, besides confirming the impact of constraints on quantification such as the Scope Constraint, also verified that at least the following factors seem to play a role:

1. a preference to assign subjects wide scope over their clausemates in active sentences;
2. a preference for NPs that denote agents of events to take wide scope;
3. a preference for definite NPs to take wide scope, and for indefinite NPs in object position to take narrow scope;
4. the distributive hierarchy observed by van Lehn.

The experiments of Kurtzman and MacDonald also give us insight in the way some of these 'principles' appear to operate. They seem to behave like defaults, in the sense that they apply unless they are overridden by stronger defaults or generate a conflict with another default with the same strength.

## 1.4 SCOPE DISAMBIGUATION IN CONTEXT: THE TRAINS DIALOGUES

An account of discourse interpretation processes such as scope disambiguation must ultimately be rooted in the analysis of natural language conversations. In this section, I supplement

<sup>43</sup> K&MD found a further set of examples problematic for this proposal besides those mentioned by Inop, namely, the examples with a quantifier embedded in a complex nominal like (1.56).




Figure 1.1 The map used by the participants in the conversation

the results presented in the previous section with an analysis of the scoping preferences that can be observed in a corpus of spoken dialogues. The data I used to study the scoping behaviour of operators in context are the conversation transcripts contained in the TRAINS corpus collected at the University of Rochester in partnership with the transcripts described in [Gross *et al.*, 1993]. Of course, studies limited to one class of conversations have to be taken with a grain of salt for this reason. I discuss here only those scoping preferences that are most clearly supported by the data and are most consistent with what is known about discourse interpretation.

#### 1.4.1 The TRAINS Corpus

The aim of the TRAINS project [Allen and Schubert, 1991, Traum *et al.*, 1994] is to study task-oriented spoken language conversations. To make things more concrete, a TRAINS domain has been introduced: the conversations we study take place between a railway manager (the *user*) whose task is to transport goods such as orange juice by train between cities of upstate New York, and an assistant (the *system*) whose task is to help the user in developing a plan. An example of the kind of task the user has to develop plans for is given in (1.60).

(1.60) I have to get one tanker of orange juice to Avon, and a boxcar of bananas to Cerning, by 'perm

The work on TRAINS is driven by the study of transcripts of recorded conversations.<sup>12</sup> The 'user' and the 'system' participating to these conversations are separated by a wall and communicate via microphones and headphones: each of them has a copy of the map in Fig. 1.1.

An example of an unedited transcript is shown in Fig. 1.2. At the current stage, the work on TRAINS is mostly based on edited transcripts such as the one in (1.61). The user's utterances are marked with 'U', the system's utterances with 'S':

<sup>12</sup> Every year a new phase of data collection takes place: the work relies on the transcripts of the conversations collected in 1993, used in [Traum *et al.*, 1993].

<sup>13</sup> Each utterance has an identification tag consisting of two numbers: the first number is the turn number, while the second indicates the different participant within the turn. This notation is due to David Traum.

(1.61)

- 13.1 U: out at the same time  
13.2 OK  
13.3 We're gonna back up engine E1 to the boxcar at Elmore.  
13.4 and send that off to Cerning.  
13.5 now while we're looking that boxcar with oranges at Cerning  
13.6 we're gonna take the engine E1  
13.7 and send it over to Cerning.  
13.8 back it up to the tanker car.  
13.9 and send it back to Elmore  
14.1 S: okay  
14.2 We could use engine E1 to take back the tanker.  
and the boxcar to Elmore  
14.3 U: oh, we can do that.  
14.4 S: yeah  
14.5 then hop the whole thing with engine E1  
14.6 and get back up the tanker car with the boxcar that has  
oranges on it.  
14.7 and take it back to Elmore  
14.8 S: okay  
14.9 that's no problem  
15.1 U: okay  
15.2 yeah  
15.3 while that's happening,  
15.4 take engine F1 to Dunville.  
15.5 pick up the tanker  
15.6 and come back to Avon  
15.7 S: okay  
15.8 U: okay  
15.9 then load the boxcar with bananas

One of the goals of the TRAINS project is the development of a natural language understanding system whose capabilities should match those of the human 'systems' that participate in these conversations. So far, four prototypes have been developed, one per year starting in 1990; the TRAINS-93 prototype will be described in some detail in Chapter 7.

#### 1.4.2 Operators in the TRAINS Corpus

The first fact to emerge from an analysis of the TRAINS corpus is the existence of a mismatch between the literature on scope and the data one can find in the corpus. While the literature on scope reviewed in §1.3 concentrates on quantifiers, very few instances of use of quantifiers other than definite and indefinite NPs can be found in these conversations. The scopal ambiguities observed in the TRAINS dialogues are originated by definite and indefinite NPs, tense, modals, negation, connectives, and wh-phrases.

The use of operators in the TRAINS conversations is summarized in the tables 1.1-1.4

- 110 Speaker: Utterance
- 11 (1) okay the problem is we better ship a buscar of oranges to Bath by 8 AM
- 21 S okay
- 31 (1) now when we need to get a buscar to Corning where are oranges at Corning right?
- 41 S right
- 51 (1) so we need an engine to move the buscar right?
- 61 S right
- 71 (1) where are engine at Avon right?
- 81 S right
- 91 (1) so we should move the engine to Avon
- 101 S engine E1
- 111 (1) E1
- 121 S okay
- 131 (1) engine E1 to Bath in time
- 141 S okay
- 151 (1) um and hook up the buscar to the engine move it from Danville to Corning
- 161 S okay
- 171 (1) how does THAT work?
- 181 S that gets us to Bath at 7 AM and (mc.)
- 191 (1) good
- 201 S ok

Figure 1.2 An unedited TRAINS transcript

the	we have to get engines to the buscars taking the engine from Avon going to Danville, pick up the buscar and bring it back to Avon
when	I have to ship a buscar of oranges to Bath by 8 o'clock today
	I see that there's a tank car there
all	we don't want a tank car
everything	while this is all happening
any	I don't even need to um ok everything will go through here let me try and make sure that everything will go through here are there buscars anywhere except at Bath?
	are there any oranges already in any of the buscars?
	any suggestions?
	is it true that we don't have any orange juice made right now

Table 1.1 Determiners found in the TRAINS transcripts, June 6, 1991

and	and hook up the buscar in the engine --
	I have to get one motor of OJ to Avon, and a buscar of bananas in Corning, uh by 3 PM and it's now 12 midnight.
before	in time enough that I can take the Avon to Corning route before the orange juice is ready to take the Corning to Avon route
if	- if I am not mistaken we need to move the orange juice in the tankcar.
	so if you were to go right to Corning, via Danville you'd get to Corning at 11
or	to pick up a buscar at Danville or Bath, or sending it, engine from Avon to pick up a buscar at Danville or Bath
	what would be faster, to send an engine from Elmore to one of the buscars or from Avon?
then	to form send the send the send the orange car and then and load the tankcar
	ok, then bag the whole thing with engine E3
	let's do it that way then
while	now while we're loading that buscar with oranges at Corning we'll go take the engine E3
until	we don't get to Avon until 4 PM

Table 1.2 Connectives found in the TRAINS transcripts, June 6, 1991

so	one buscar of oranges makes one tankcar of orange juice so we need send 8 AM, we send E3 to Corning and then to Danville to pick up the buscar
um	oh, we don't want a tankcar, do we?
	why don't we pick the one up at Danville, that seems to be the next closest one.
where	we don't get to Avon until 4 PM.
and	and where does that have to go... to Avon.
just	okay I'll right I misunderstood.
did	in fact, what did I say, did I say oh, did I send a buscar from Bath to Corning?
didn't	yes all right I didn't see, but buscar
before	even let's see if the buscar problem is going to conflict with any of this
can	okay, I can make a whole plan at this point right?
can't	you can't have two trains running on the same track.
could	um, we could do that, but the problem is
couldn't	we couldn't send it
have to	we have to get one tankcar car of orange juice to Avon
should	engines should be shipped in a buscar
	clearly we should go Avon to Corning via Bath
would	and which mode do you want to take to Avon?
	what would be faster to send an engine from Elmore to one of the buscars or from Avon?

Table 1.3 Tense, auxiliaries and negation in the TRAINS transcripts, June 6, 1991

what	... so, I think what we should do is hook up uh one of the engines at Elmore
when	... what time do we get to Danville?
where	... what if we went to Avon how long is it, Avon ... when the (C) is read, you head it up into the tanker car and bring it back to Corning
where	... when we do have range, you call I mean?
what	... so we need to get a hauler in Corning, where there are oranges. If and then where do we have to get with that?
how	... where are the oranges?
	... which route do you want to take?
	... which engine are you want?
	... which is the fastest route to Avon?
	... instead of sending along F1 to Danville, which is where we go / had it before, ...
	... how long is it from Corning to Danville, where there's that other hauler?

Table 1.4: Wh-phrases in the TRAINS transcripts, June 6, 1991

#### 1.4.3 The Scope of Definite Descriptions

The determiner "the" is the fourth most common word in our dialogues and by far the most important disambiguating factor one can observe in the TRAINS corpus is the tendency of definite descriptions to take wide scope, which is consistent with van Lier's observations. A clear contrast between the scoping of definites and the scoping of indefinites can be observed:

- (1.62) a. We need to send an engine to Corning.  
b. We need to send the engine to Corning.

In (1.63) and (1.64), the definite NPs "the warehouse" and "the engine" take wide scope over the modals "need" and "should", respectively:

- (1.63) we need to send uh a hauler to the warehouse and an engine to the warehouse

- (1.64) we should get the engine to pick up the hauler and haul for Corning.

In (1.65), the definite takes scope over a wh phrase:

- (1.65) Where are the haulers?

#### 1.4.4 The Scope of Modals

Equally clear is the preference for modals not to take wide scope with respect to definites, but wide scope with respect to indefinites. A typical example is (1.66):

- (1.66) We should hook up an engine to the hauler at Avon.

The preference for modals to take wide scope over indefinites is not affected by the syntactic position of the indefinite, as shown by (1.67). (It should be noted that a virtually no indefinite in our dialogues is specific.)

- (1.67) An engine should get to Corning to pick up the hauler.

An interesting interaction with purpose clauses is shown in (1.68):

- (1.68) We should send an engine to Corning to get the hauler.

The adverbial may either express the purpose of the action of sending an engine to Corning, or the purpose of the necessity statement.

#### 1.4.5 Syntactic and Semantic Constraints

At least two kinds of constraints operate in the TRAINS dialogues. First of all, there is a Scope Constraint effect preventing a modal appearing in one conjunct of a coordinated structure to take scope over the other conjunct. An example is sentence (1.69), that does not have a reading in which the modal in the first sentence, "can," takes scope over the whole coordinated structure.

- (1.69) We can go to Danville back to Corning and that's perfectly compatible with our other little plan.

We also find instances of a constraint observed by Milsark (Milsark, 1971) reported in Heim, (1987), who observed that indefinites in post copular position in three insertion sentences, cannot take scope outside their clause.

- (1.70) There is a book for everyone.

- (1.71) John believes that a thief is in his apartment.

- (1.72) John believes that there is a thief in his apartment.

In our dialogues we can find examples as in (1.73):

- (1.73) a. I believe that there is a tanker car already there.

- b. I think there'll be no problems with that.

- c. There isn't a hauler at Avon.

#### 1.4.6 Further Observations

In conclusion, I mention a fact discovered while examining the corpus that, although not addressed in this dissertation, has nevertheless affected some of the technical solutions. Consider the fragment in (1.74). The modal "could" in 13.2 'takes scope' over the indefinite "some oranges" in 15.3, that is part of a separate utterance. Conversely, the indefinite "the engine" in 15.1 'takes scope' over the modal. This indicates that scope interpretation in our dialogues has to be addressed at a level in which a distinction is made between sentences and the utterance events that, once combined, result in these sentences. The model of interpretation I propose in the dissertation has been developed with this kind of problems in mind.

- 9.1 U: so we should move the engine at Aveno.  
 9.2 : engine E, in... (on)  
 10.1 S: engine E  
 11.1 U: E!  
 12.1 S: okay  
 13.1 M: engine E1, to Bath, to (ec)  
 13.2 : re, we could actually move it to Danville, and pick up  
       : the buscar there  
 14.1 S: okay  
 15.1 U: um and back up the buscar in the engine.  
 15.2 : move it from Danville to Coving.  
 15.3 : find up some cranes into the buscar.  
 15.4 : and then move it on to Bath.  
 16.1 S: okay.

(1.74)

## 1.5 OVERVIEW OF THE PROPOSAL

### 1.5.1 Recapitulation

I began this chapter by trying to answer a question that, surprisingly enough, is not usually addressed in work on discourse interpretation: what does the Combinatorial Exclusion Paradigm tell us about the way humans go about interpreting utterances? The answer I gave was centered around the distinction between *semantic ambiguity* and *perceived ambiguity*, and consists of two hypotheses about the way humans do discourse interpretation: the *Underspecification Hypothesis*—humans are able to represent semantic ambiguity implicitly—and the *Anti-Random Hypothesis*—although humans usually attempt to resolve ambiguity, they do not do so by randomly generating alternative interpretations or “logical forays”; those interpretations are only generated that are suggested by context or by discourse interpretation procedures.

Having thus clarified my assumptions about discourse interpretation, I turned to the specific discourse interpretation process I am concerned with, scope disambiguation. Reviewing the literature on scope disambiguation, I noted that, while a certain number of facts about scope disambiguation have been observed with some consistency—the important role played by constraints, the preference for subjects in active sentences to take wide scope, the tendency of certain quantifiers to take wider scope than others, and in particular the tendency of definites to take wide scope—the picture of scope disambiguation that emerges from this work is far from coherent. What is known is formulated as a set of stipulated and largely unrelated “principles,” whose interaction hasn’t been thoroughly studied. This lack of understanding is reflected by the way scope disambiguation is done in NLP systems, where the scope disambiguation principles proposed in the literature are used as scoring factors; both the choice of some principles rather than others, and their relative weight, are determined by trial and error.

I then argued that further insights about the nature of disambiguation principles, and especially about their interaction, can be obtained from Kuntzman and MacDonald’s recent experiment, and by looking at the scope preferences one observes in actual conversations. The results of Kuntzman and MacDonald suggest that whatever the factors are that play a role in scope disambiguation, it is plausible to assume that they are always “active” (e.g., it’s not the case that sometimes the principle that results in subjects taking wide scope operates, whereas at other times it doesn’t), but they behave like prioritized defaults, in the sense that they can be overridden by defaults with higher priority, and originate conflicts with defaults with the same priority.

Data about scope preferences in naturally occurring conversations was obtained by examining the transcripts of the TRANS corpus. Besides confirming the importance of constraints and the tendency of definites to take wide scope observed in the literature, my analysis showed that modals tend to take wide scope over indefinites.


### 1.5.2 Discourse Interpretation as Inference over Underspecified Representations and the DRS Construction Algorithm

If the Underspecification Hypothesis and the Anti-Random Hypothesis are correct, then the inferences that humans make during discourse processing must not require a preliminary step in

which an unambiguous representation is obtained by random guessing. Some of these inferences must rely only on the kind of information provided by the specified representations.

In other words, discourse interpretation must have some of the characteristics of the DRS construction algorithm proposed in Discourse Representation Theory (DRT), a theory of discourse interpretation developed by Kamp and Reyle (1981, Kamp and Reyle, 1993). The goal of the algorithm is to construct a Discourse Representation Structure (DRS) out of sentences.<sup>14</sup> According to this theory of discourse processing, discourse interpretation begins when the s-structure of the sentence whose interpretation has to be compared is added to the common ground, this new element of the common ground is called an *uninterpreted condition*. Model Construction Rules (this term is mine) can then apply, these are rewriting rules that replace the uninterpreted condition with one that is less ambiguous.


Consider for example the sentence "Every kid climbed a tree". According to the DRT construction algorithm, the first step of interpretation is to add to the common ground the underspecified representation of the sentence, that in DRT is its s-structure. The common ground itself is a Discourse Representation Structure, an object that consists of a set of discourse markers and a set of conditions represented as a 'box'. The result of this initial step of interpretation is shown in (1.75).



(1.75)

The DRS in (1.75) is the preliminary 'partial hypothesis' about the sentence. Now the model construction rules may apply, in this case we have a choice between applying the rule for indefinites or that for universally quantified NPs. Both of these rules are 'triggered' by the occurrence of a certain syntactic pattern in a DRS, and perform two tasks: adding some material to the common ground, and/or rewriting the current partial hypothesis. Say that the rule for indefinites applies first, this rule adds a new discourse marker to the DRS containing the pattern it's triggered the rule, and rewrites the pattern so that the partially disambiguated representation in (1.76) is obtained. This new DRS consists of a new discourse marker, *x*, and two conditions: one to the effect that *x* has the property of being a tree, while the other is an uninterpreted condition obtained by replacing the top *NP* of the tree with the new marker *x*.

<sup>14</sup> The procedure for building a discourse model is based on the version of DRT developed by Heim (1982) as follows: interpreting the purpose of the current discourse.




(1.76)

DRT provides an elegant example of underspecified representation and *c'* operations on underspecified representations. The format used for underspecified representations is used to capture syntactic constraints and a representation for 'partially disambiguated hypotheses' (DRSs containing an uninterpreted condition) is also given. Kamp and Reyle's book contains the most detailed analysis currently available of the operations on underspecified representations that lead to a disambiguated interpretation.

The discourse interpretation procedure developed in DRT does not take into account the problem of ambiguity in natural language, which is not surprising, as that was not one of its goals. No provision is made in DRT to account for the difference between semantic and perceived ambiguity, and indeed there is no way to capture perceived ambiguity. No constraints on disambiguation are specified in DRT, the DRS constructor *Construct* is supposed to generate all of the alternative readings of a sentence, therefore it is intentionally designed in such a way as to suffer from combinatorial explosion. Uninterpreted conditions have in DRT the same status as underspecified representations had in early NLP systems: they are treated as uninterpreted representations and do not reflect a commitment to any particular theory of semantic ambiguity.

In other words, one can't just take DRT as it is and get a theory of surface discourse interpretation that gives a satisfactory treatment of ambiguity. Nevertheless, as I'll show in the rest of the dissertation, the standard version of the algorithm can be modified to yield a theory of discourse interpretation consistent with the hypotheses discussed in this chapter. The rule format developed by Kamp and Reyle to describe discourse processes such as tense interpretation can be used to formalize the other aspects of discourse interpretation, such as scope disambiguation. I will also show that DRT's 'uninterpreted conditions' can be replaced by underspecified representations that do represent a commitment to a particular theory of semantic ambiguity and for which an interpretation can be provided, while at the same time maintaining the basic form of the algorithm. I will call these representations 'logical forms'. Logical forms codify both information about the s-structure of an utterance and information about the semantics of its elements; they can be represented as s-structures whose leaves consist of the semantic translation of lexical items and/or functional categories. The logical form for the sentence 'I can find a piece of paper' is shown in (1.77).





I'll conclude by adding that my decision of using DRSSs as representations also has a pedagogical motivation: namely, the desire to give a graphical representation to my intuition that the scope preferences observed in the literature are not obtained by choosing between permutations of operators at logical form, but rather by building descriptions of situation types and establishing dependency relations among them.<sup>16</sup> I believe that the main reason why we don't have a good theory of scope disambiguation is because, so far, people working on scope interpretation have made the mistake of focusing on scope as a structural relation between operators in logical forms, instead of looking at what the semantic notion of scope is an abstraction of. In part because of the influence of work on scope in the generative tradition, the emphasis has been on methods of scope interpretation based on logical form manipulation, which almost unavoidably led to algorithms suffering from combinatorial explosion problems.

### 1.5.3 Contextual Interpretation and Scope

Let us now return to the scoping preferences observed in the TRAINS corpus. Should we attempt to account for those by stipulating a couple of additional disambiguation principles? I believe that no new principles are necessary. In fact, I believe that the facts observed in the TRAINS conversations may be the basis for a theory of scope disambiguation that reduces the 'principles' observed in the literature to facts about discourse interpretation.

First of all, I'll note an interesting correlation that can be observed in these conversations. The operators for which clear scoping preferences exist (definites and modals) share two important properties: (i) the meaning of both definites and modals is in part determined by context, and (ii) in the TRAINS dialogues, appropriate values for these context-dependent aspects of interpretation are very salient. On the other hand, the NPs which tend to take narrow scope (indefinites) do not tend to receive context-dependent interpretations such as specific interpretations.

That definites are context-dependent is generally accepted. In fact, this is taken by many (e.g., Heim [1982] or Hawkins [1978]) to be their defining property. Heim defines definites in the TRAINS conversations as either anaphoric or referring to the 'visual situation' shared by the participants in our conversations, namely, the map (see §1.4). These kinds of references can be interpreted easily, and anyway we know from work by Craun and Steedman and Altmann and Steedman that definites are usually interpreted early, in early in fact that their interpretation may even affect parsing [Craun and Steedman, 1985; Altmann and Steedman, 1988].

While the hypothesis formulated by Kratzer [Kratzer, 1977; Kratzer, 1981], that the meaning of modals is context-dependent is not as universally accepted as the hypothesis that definites

<sup>16</sup> That the DRSSs form situations is a good way to represent situation types can be seen from the recent work on the Extended Kamp Notation for Situation Theory [Barwise and Cooper, 1993].

are, this hypothesis nevertheless provides the most reasonable explanation of the fact that modal expressions such as "must" may have a potentially infinite number of interpretations, of which the traditional deontic and epistemic interpretations are only two examples. (1.78a) and (1.78b) are examples of a deontic and epistemic reading of a modal. In (1.78a) the modal "must" is interpreted *deontically*—the interpretation of the sentence is something like "In all worlds that behave according to the customs of the Maori, all children learn the name of their ancestors"—while the interpretation of "must" in (1.78b) is *epistemic*: the sentence means something like "In all worlds that are consistent with what we know, the ancestors of the Maori arrived from Tahiti". Kratzer notes, however, that additional interpretations are available: (1.78c), for example, although also interpreted deontically, now refers to the US customs: that sentence can be paraphrased as "In all worlds that behave according to the US customs, the children go to school until eight grade."

- (1.78) a. All Maori children must learn the names of their ancestors  
 b. The ancestors of the Maori must have arrived from Tahiti  
 c. All US children must go to school until eight grade

Kratzer argues that modals are not ambiguous, but have a context-dependent meaning. Much as when interpreting a quantifier the domain of interpretation has to be contextually restricted so that "Everyone is asleep" does not literally mean that every human being sleeps, when interpreting a modal a listener has to identify its modal base—the set of worlds over which that modal quantifies. In the case of (1.78a), for example, the modal base is the set of situations consistent with the customs of Maori.

As it happens, in the TRAINS conversations the modal base with respect to which modals have to be interpreted is always very clear: it's the set of worlds that are consistent with the plan being elaborated by user and system. This modal base is so salient, in fact, that virtually every modal in the corpus is interpreted with respect to the plan.

So, two of the operators that tend to take wide scope in the TRAINS conversations have a context-dependent aspect that can be resolved very easily in these conversations. By contrast, we can see that indefinites, that usually take narrow scope in the TRAINS conversations, are never interpreted specifically, because in the TRAINS domain boxcars, engines and tanker cars are all equivalent, so it's highly unlikely that a listener will ever suspect that a speaker, when using the indefinite NP "a boxcar," refers to a specific boxcar.


It is natural, then, to advance the hypothesis that 'scope assignment' depends on discourse interpretation, at least as far as definites, indefinites and modals are concerned. In a DRT-like framework, we can formalize this claim as the hypothesis that whether or not the modal construction rule for an operator applies depends on the results of discourse interpretation.

**Condition on Scope Disambiguation (CSD)** : The resolution of the contextual component of an operator's meaning takes place before its 'scope' has been determined. The model construction rules that specify the contribution of an operator to the common ground apply if, and only if, the contextual aspects of an operator's meaning have been determined.

As this idea plays a central role in the dissertation, it is worth illustrating it with a couple of examples. Consider first of all the case of modals taking scope over indefinites, illustrated by (1.79):


## (1.79) An engine must go to Avon

According to Kratzer's theory, this sentence, prior to the determination of the modal base, has the following truth conditions: the sentence is true in a situation if all situations that are part of the modal base determined by that situation can be extended into situations in which an engine goes to Avon. The underspecified representation/lexical form for this sentence is shown in (1.80).<sup>36</sup>




The interesting aspect of the representation in (1.80) is that the translation of the lexical item "must" includes a parameter  $k$ . The presence of a parameter in the interpretation of a lexical item indicates that some aspect of that interpretation hasn't been resolved yet: in this case, what still has to be determined is the modal base of "must," translated in (1.80) as a function from properties  $P$  to properties true of an object  $x$  iff all situations that are an instance of the situation type  $k$  (yet to be determined) include a subsituation  $s$  that is an instance of the kind of situation in which  $x$  has the property  $P$ . The indexical constant  $g$  refers to the situation at which an expression is evaluated. (The representation is presented in more detail in Chapter 4, while the translation of modals is discussed in detail in Chapter 5.)

This logical form is the starting point for discourse interpretation. Because the plan (a situation kind, as discussed extensively in §6.1.2) is very salient in the TRAINS conversations, the hypothesis immediately suggests itself that the plan should serve as the modal base for the modal "must." Once this assumption is made, the model construction rule for modals is licensed, this rule, much as the rule proposed in standard DRT, results in the partially disambiguated interpretation in (1.81), where the modal takes scope over the indefinite, that has not been interpreted. In fact, no further disambiguation is needed, because the logical form in (1.81) is equivalent to the interpretation that would be obtained if the existential were to be rewritten in the standard DRT fashion. The representation in (1.81) can be paraphrased as follows: all situations that are an instance of the plan contain a subsituation  $s$  such that an engine goes to Avon in this situation. This is the interpretation consistent with the preferred scope




<sup>36</sup> The analysis of indefinite NPs and modals in terms of existential quantification will be replaced by one based on DRSs in Chapter 5.

In the case of definites, I assume, following Hawkins (Hawkins, 1978), that the contextually dependent aspect of their meaning is their resource situation—the set of objects in the common ground from which the referent of the definite description is selected. The logical form of the sentence "I don't see the engine" is shown in (1.82); the translation of the definite "the engine" is the set of properties that hold of the denotation of  $x$ , an object that is an engine in the resource situation  $s$ , yet to be determined.<sup>37</sup>



The discourse processes that interpret definite descriptions, just as the processes for interpreting modals, operate on the basis of the current context and the information provided by this logical form and, in the case of a felicitous use of the definite description, result in an hypothesis about the resource situation.<sup>38</sup> Once the contextually dependent aspect of the interpretation of a definite has been resolved, the appropriate model construction rule may apply: the result is a partially disambiguated interpretation of the form in (1.83). This DRS is the scriptally preferred interpretation: if "I don't see the engine."



The point of these examples is to show that the facts about the scope of definites and modals in the TRAINS conversations can be explained without stipulating a principle stating that modals take wide scope over indefinites, and one saying that definites take scope over negation. In fact, I will show in Chapter 6 that, once analyses of definite description interpretation, modal interpretation, and tense interpretation are provided, we can get the most important scoping preferences observed in the TRAINS conversations without the need of stipulating any such principles. The only other factors that seem to play a role, besides the Condition on Scope Disambiguation seen above, are constraints on model construction rules preventing the generation of interpretations that violate the definiteness restriction and the island constraints.

<sup>37</sup> The representation assumed here for definites is the one used in DRT, where definites and indefinites introduce free variables—discourse markers—on the representation. In (1.82),  $x$  is a discourse marker.

<sup>38</sup> I am assuming here the theory of definite description interpretation proposed in (Poesio, 1993) described in §6.5.

### 1.5.4 The Results of Kurtzman and MacDonald, Revised

Does the hypothesis that scope disambiguation is driven by discourse interpretation processes that have nothing to do with 'scope' hold when operators other than definites, indefinites and modals are concerned? In particular, what happens with quantifiers? Where do the preferences observed in the literature come from?

I propose that the scope preferences discussed in §1.3, as well, can be explained as the result of discourse interpretation processes. I shall concentrate on the data discussed by Kurtzman and MacDonald. As you may recall, they observed that subjects tend to take wide scope in active sentences, especially if the verb requires an agent, and that in passives no preference can be observed, concluding that at least two different factors were at play. I propose that these preferences, as well as the existence of a conflict, can be made to derive from what we already know about discourse interpretation. I am referring in particular to the following two hypotheses about discourse interpretation

- A person, upon reading or hearing a sentence, tries to identify the part of that sentence that is given and the part that is new, in the absence of an already established context; this partitioning process is driven by information provided by commonsense knowledge and by the syntactic structure of the sentence (Clark and Haviland, 1977)
- A person is not only trying to establish a connection between the last sentence and the previous context, he or she is also concerned with building a description of the situation described or desired. This is done by associating predicates with their arguments, a process called *argument selection* in which 'thematic' information plays an important role (Jackendoff 1972; Dowty, 1989; Dowty, 1991; Grimshaw, 1990)

It should come as no surprise to those who believe in Grice's Relevance maxim (Grice, 1967) that well-constructed utterances should consist of a part that establishes a connection with the existing discourse, and of a part that provides new information (Clark and Haviland, 1977). Langue, as provides tools for indicating what's given and what's new in a sentence: for example, a question like "Who broke the window", that creates a context in which the occurrence of an event of breaking the window is established, and in which new information is explicitly requested, can be felicitously answered by means of (1.84b) or (1.84d), but not by using (1.84c) or (1.84e).

- (1.84) a A Who broke the window?  
b B It is John who broke the window  
c B #It is the window that John broke.  
d B JOHN broke the window  
e B #John broke THE WINDOW.

The given/new structure of a sentence can be explicitly indicated by using *it*-cleft clauses of the form "It is ... that ...", such as (1.84b), where the material in the *it*-cleft clause is new, or by using focal stress as in (1.84d), where the stressed elements of the sentence (in capital letters) are again interpreted as new information, or by using fronted PPs, where the material in the fronted PP is taken to refer to given or background information relative to the rest of the


sentence (Clark and Haviland, 1977; Prince, 1981; Delin, 1989). A weaker indication of the way a sentence is split in given and new information is provided by the subject/predicate split: subjects are typically taken to indicate given information, whereas the rest of the sentence is usually taken to provide new information (Clark and Haviland, 1977)

We can arrive at the relative between the given/new partition and scope if we keep in mind that the semantic notion of 'scope' is really meant to capture *semantic dependency*: if item A is 'in the scope' of item B, then the interpretation of A is (or may be) affected by the interpretation of B. A depends on B. On the other hand, the interpretation of new information depends on the rest of the context, that is, on what's given: we predict therefore that the part of a sentence's information that is inferred to be 'given' will 'take scope over' (will be independent from) the part that's 'new'.

We have already seen an example of construction that supports this prediction; definites, definites are 'given by construction' and therefore their interpretation is usually independent from the interpretation of the rest of the sentence. Of the other constructions that explicitly mark the given/new partition in a sentence, one (*it*-cleft constructions) doesn't allow the use of quantifiers in the cleft clause, presumably because of scope island constraints (compare "It is every kid that climbed a tree"). The effects of stress on scope are also difficult to quantify because of the overwhelming preference for interpreting stress contrastively in these constructions (compare "Every KID climbed a tree"). Even in the case of proposed PPs and in the case of subjects, however, the prediction that what's given will 'take wide scope' over the rest of the sentence correlates well with the judgments found in the literature. (We have discussed the data about subjects extensively in §1.3.) The clearest examples of the effect of preposing on scope are those due to Reinhart (1997) discussed above. In (1.85a), both the reading in which 'some reporters' takes wide scope and the reading in which 'every town' does are available, although the first one is generally preferred; in (1.85b), only the second reading is available.


- (1.85) a Some reporters followed Kissinger in every town.  
b In every town, some reporters followed Kissinger

The effect of the given/new contrast on scope can be understood more clearly once we spell out more in detail what it means for part of a sentence to be given information. This issue is discussed more in detail in Chapter 6, but it can be summarized as follows. Consider again the sentence "Every kid climbed a tree", whose initial underspecified interpretation is shown in (1.86)





A listener confronted with this sentence, and with no other information about what's in the context, will have to make a more or less educated guess. On the basis of his or her knowledge about the way English sentences are structured, the listener is likely to conclude that the appropriate interpretation is one where the context includes a set of kids, and a property is predicated of these kids. In other words, the listener will add to his/her perspective on 'common ground' a statement to the effect that the speaker/writer is talking about a situation

which contains a set of kids, and this situation serves as resource situation for the quantifier "every kid". What gets added to the common ground may be schematically represented as in (187), that may be paraphrased as follows: the existence of a situation  $x$  in D is hypothesized, about which we have the information that it contains a set  $X$  of kids (the initiative  $T$  to indicate predicates that denote sets of objects of type  $T$  is due to Link (1981)). This is the situation that the speaker is referring to when talking about "every kid".



Now the model construction rule for universal quantifiers is licensed, the result is the partially disambiguated hypothesis in (188).




Before proceeding to the next discourse interpretation procedure, let me remark that the process just discussed is essentially a process of *presupposition accommodation*, closely related to the proposals of Heim (1983a) and van der Sandt (1990). In general, it is one of the claims of this dissertation that some of the scope preferences observed in the literature may express judgments about dependency or independence of elements of a sentence's meaning that are best captured not in terms of scope, but of other forms of dependency, such as presupposition.

Coming to the other discourse interpretation process I mentioned above, *argument selection*, I will first of all note that two hypotheses are usually made in the literature discussing the process by which the syntactic structures expressing a certain predication relation are generated. The first hypothesis is that the notion of *thematic role* (Gruher, 1965; Fillmore 1968; Jackendoff, 1972; Jackendoff, 1983; Dowty, 1989; Dowty, 1991) plays an important role in this process. The second hypothesis is that certain roles are 'more equal than others', not only are there rules for mapping thematic roles into syntactic constituents or viceversa, there is an *order* in which this mapping occurs. This second hypothesis, motivated by data about passivization, reflexives, and the theory of control (see, for example, Jackendoff, 1972; Nishigauchi, 1984; Dowty, 1991) comes in various versions. Some believe in the existence of a *thematic hierarchy* (e.g., Jackendoff, 1972; Grimshaw, 1990), while others simply propose that agents are 'distinguished arguments' (Koenig, 1993).

Whichever version of the theory is chosen, the hypothesis that the situation described by a sentence is constructed beginning with agents, if the verb is agentive and a proper filler can

be found besides giving us independent motivation for the preference for subjects of active sentences to take wide scope when the main predicate requires an agent, also provides us with an explanation both of the contrast between perception verbs and action verbs observed by Karttunen and MacDonald, and of the contrast between actives and passives, since in the case of passives, the interpretations obtained by the process that identifies the given/new partition and assigns arguments to predicate conflict.

The interpretation obtained for the sentence "Every kid climbed a tree" by the argument selection procedure is much the same as that obtained by the procedures that implement the given/new contrast: the listener tries to identify the agent(s) of the sentence he or she hears or reads, if no contrast is provided, the subject obtains an interpretation by adding to the common ground a new set of objects of the appropriate type, and by establishing a relation between this set of objects and the resource situation of the NP in agent position.<sup>15</sup>



### 1.5.5 Interactions Between the Discourse Interpretation Procedures That Affect Scope and the Architecture of Discourse Interpretation

Karttunen and MacDonald suggest a *parallel* architecture for discourse interpretation of the type displayed in Fig. 1.3, where all procedures for discourse interpretation (at least, all those at the same level, see below) simultaneously work off the initial underspecified representation I also observed, discussing Karttunen and MacDonald's experiments, that their results seem to indicate that the processes involved in scope assignment appear to behave like defaults. The architecture in Fig. 1.3 is consistent with the perspective on reasoning by default found for example, in Reiter's Default Logic (Reiter, 1980) according to this view, defaults behave like forward-chaining rules. Each hypothesis obtained by discourse inference corresponds to what

<sup>15</sup> The one factor that Karttunen and MacDonald did not test in their experiments is the Quantifier Hierarchy. Their experiments, however, do seem to support van Loh's thesis that there are two distinct hierarchies, rather than (as I proposed) that only one quantifier hierarchy exists, as shown, for example, by their data about complex NPs. Of the two hierarchies proposed by van Loh, the Specific/Inherent hierarchy was introduced to explain preferences that already follow from the model presented here without further stipulations. As far as the data that motivated the Quantifier Hierarchy are concerned, they are in part the result of the fact that NPs such as "each house" or "every house" semantically denote quantifiers (in DRT terminology, they introduce ternary operators on the common ground), while others, such as plural definites, denote sets or groups (the semantic translation of lexical items is presented in Chapter 5) therefore do not generate dependencies when added to the common ground. The distinction between each-NPs and every-NPs is more difficult to account for unless we have a good account of each-NPs, but the data are also much less clear so I will leave that issue aside in this dissertation.






Figure 1.4 The interaction of strong and weak defaults

them as more plausible. Examples of use of commonsense knowledge for filtering are shown in (1.91) in both examples, the choice of an interpretation depends on the semantic properties of the concepts involved ('tutoring' and 'assigning', 'workstation' and 'sitting').

(1.91) a A tutor is assigned to every student

b A workstation sits in every office

The preferred interpretation of (1.92) is the one where "it" refers to the serial port, because these are usually found in the back of computers. This commonsense knowledge-cumulated 'plausibility ranking' among hypotheses can be formally modeled by means of a priority ranking among models similar to that proposed by Shoham (Shoham, 1988)

(1.92) Hook up the cable to the serial port. It is on the back of the computer

Commonsense reasoning may also tell us that some extensions are equivalent for the purposes at hand, thus can be merged. An example from the TRAINS corpus is shown in (1.93): even though "it" can refer either to 'the engine' or to 'the boxcar', and therefore two extensions could be obtained by discourse interpretation, the difference between these two extensions would be immaterial as far as the plan is concerned, because moving one object would necessarily entail moving the other: the two hypotheses can therefore be merged

(1.93) Hook up the engine to the boxcar, and move it to Avon

The only situations in which an ambiguity is perceived are those in which commonsense knowledge cannot rank or merge the hypotheses obtained by discourse interpretation. What happens at this point depends to some extent on the strength of the information used to arrive at the interpretation. If weak defaults were used, as in the case of passives, the conflict might be resolved by picking up one interpretation at random, as in Kurtzman and MacDonald's

experiments. If the conflict is between two reliable sources of information, then the conflict cannot be resolved and we have a true perceived ambiguity

Besides its role in ranking and filtering hypotheses, commonsense reasoning (I am including inferences based on world knowledge here) plays at least two other distinct roles in discourse interpretation. One of these roles is to suggest certain interpretations. Consider the following example:

(1.94) Every school sent the principal to the meeting

The preferred interpretation of this sentence is the one in which 'the principal' is in the scope of 'every school'.<sup>1</sup> The intuitive explanation for this preference is that the predicate 'principal' is primed or 'in implicit focus' in a context in which the predicate 'school' has been used, and therefore the preferred interpretation is the one in which an association between the two NPs is established. Or consider again, instead, van Lehn's example

(1.95) Every researcher at Xerox PARC knows a Lisp dialect

For someone familiar with the work done at Xerox PARC in the Seventies, the preferred reading of this sentence is the one in which the indefinite 'a Lisp dialect' is interpreted as specific, and referring to Interlisp. Again, the intuitive explanation for this preference is that the concept of Interlisp is primed in a context in which the marker 'Xerox PARC' has been introduced, and therefore the indefinite 'a Lisp dialect' can be linked to some object that is implicit in the context. That priming plays a role in interpretation has been established fairly securely in the psychological literature (Swinney, 1979; Tannenhaus *et al.*, 1979; Hirst, 1987), and there are several proposals in the artificial intelligence literature discussing how to implement priming effects and use them for interpretation—be it in terms of spreading activation (Charniak, 1986; Hirst, 1987; Norvig, 1987) or, more recently, in terms of statistical association (Hindle and Rooth, 1993)

The third role of commonsense knowledge is to complete the hypotheses obtained by discourse interpretation. Knowledge about the domain must almost certainly include causal rules that specify the results of certain actions; these rules are used to complete an interpretation (in the sense of Schank and Abelson (Schank and Abelson, 1977)), as well as to understand how the described actions 'fit together' in a plan (as discussed, say, by Kautz (Kautz, 1987)). In the TRAINS domain, for example, we know that the action of unloading the cargo of a boxcar makes that boxcar empty, after the action is performed, such a boxcar may be referred to as 'the empty boxcar'. Default assumptions are also used to 'fill in' the details of the plan left unspecified by the conversation (e.g., make sure the engine has enough gas to run, go slowly to avoid crashes, and so forth)

Even though a number of researchers may accept what I just said about the role of commonsense knowledge in discourse interpretation, currently no generally accepted theory exists of what it means for an hypothesis to be more plausible; the impact of lexical priming has not been

<sup>1</sup> English speakers tend to prefer to replace 'the principal' with 'its principal' in (1.94) but the opposite is true of French, Italian and German speakers. In Italian, 'Ogni scuola ha mandato il suo preside all'incontro', where the closest translation of 'its principal' is used, is not as good as 'Ogni scuola ha mandato il preside all'incontro', where the translation of 'the principal' occurs.

completely clarified, and finally, there are plenty of problems in formalizing world knowledge about actions and their effects so as to avoid paradoxes such as the *frame problem* (Hayes, 1987).

Providing such theories would involve writing at least one additional dissertation. The approach I will follow here is to make an hypothesis about the interface of discourse interpretation with commonsense reasoning, and leave 'placeholders' to be filled in with appropriate theories.

First of all, I propose that the only use of commonsense knowledge during the discourse interpretation phase described in the previous section is to trigger hypothesis generation by means of priming patterns. My assumption about commonsense knowledge in this phase is that a predicate *P* may be lexically primed by a predicate *Q*. I represent these associations by means of a relation *PRIME* (*P*, *Q*). By extension, I say that a discourse marker *d* is primed by the discourse marker *d'* if *P* is predicated of *d*, *Q* is predicated of *d'*, and *PRIME* (*P*, *Q*).<sup>42</sup>

Some preliminary support to the idea that commonsense knowledge is used in discourse interpretation to filter interpretations rather than suggesting them can be gathered by the examples shown above, e.g. (1.91). Note that in the case of (1.91a) the passive form was chosen, that gives less indication of the preferred interpretation, and in the case of (1.91b) the 'funny' interpretation in which the same workstation sits in every office is very much available even though commonsense knowledge would suggest otherwise. This intuition is supported by Hudson and Tanenhaus's work on centering (Hudson, 1988). Hudson and Tanenhaus looked at examples of the form of (1.96). Even though in (1.96c) the only plausible shooter is Bill who currently commits the ball, Hudson and Tanenhaus found that a significant number of subjects thought that John was the shooter, and those that 'correctly' indicated Bill as the shooter took a relatively long time to come out with the interpretation. This suggests, according to them, that a preliminary hypothesis is obtained in these cases, the hypothesis is, however, judged implausible, at which point another hypothesis has to be generated. If commonsense reasoning were used to suggest hypotheses, all subjects should come out with the same interpretation, or at least take the same time in arriving at it.

- (1.96) a John received the ball  
b He passed it to Bill.  
c He shot

Once one or more hypotheses have been obtained, commonsense knowledge is used to filter and rank them. This role of commonsense knowledge is hidden in an *evaluation function* that is part of the system of discourse inference described in Chapter 4. In TRAINS-93, this task is performed by the dialogue manager, with the assistance of the plan reasoner.

Finally the resulting hypothesis is augmented by means of inferences based on commonsense knowledge. Again, this task in TRAINS-93 is performed by the plan reasoner. While there are plenty of problems in formalizing this aspect of world knowledge in general, usable domain theories can be developed with the available formalisms, and reasonable hypotheses about the interaction of this knowledge with discourse interpretation can be made (see (Ferguson and Allen, 1993)).

<sup>42</sup> In TRAINS-93 priming is implemented in an extremely ad-hoc way.

### 1.5.7 Reflexive and Abductive Architectures for Discourse Interpretation

Another view of discourse interpretation has also been proposed, according to which discourse interpretation is an *abductive* process (Hobbs *et al.*, 1990). According to the abductive theory of discourse interpretation, a listener, when confronted with an utterance like "They went to see the movie," sets up to solve a set of interpretive problems, such as what explanation can we find for the use of the pronoun "they"? I.e., what was the intention of the speaker in using that pronoun?

While abduction may ultimately be equivalent to non-monotonic deduction, I have adopted here the latter way of formalizing discourse because I believe it better captures the hypothesis, backed by linguistic and psycholinguistic evidence, that discourse interpretation is, at first at least, a 'reflexive' process, in the sense that the first interpretive steps are largely automatic, and it's only when these initial processes do not lead to a satisfactory interpretation that an attempt is made to consider all plausible alternatives (this second phase of interpretation is sometimes called 'puzzle mode'). Support for this view of discourse interpretation is provided, for instance, by examples such as those in (1.97) and (1.98).

- (1.97) John dropped ten marbles on the floor, but he only found nine of them. ??It must be behind the couch.

- (1.98) a The car went by roaring. The radiator cap was shining  
b The car went by roaring. ??The dog barked

The use of the pronoun "it" in (1.97) is not felicitous, even though a referent for the pronoun can easily be found by reasoning about the scene. Similarly in (1.98b) the referent "the dog" cannot be interpreted as referring to a dog in the car, although it might be more likely to see dogs in cars than to see cars with radiators.<sup>43</sup> That these examples of reference are not felicitous cannot be easily explained if really discourse interpretation is simply a matter of building plausible interpretations.

### 1.5.8 SAD-93

The ideas about scope disambiguation and its relation with discourse interpretation proposed in this dissertation have been implemented in a system called SAD-93.<sup>44</sup> SAD-93 is a (*surface*) *discourse interpretation* system given the result of syntactic interpretation and lexical disambiguation as its input, it generates a set of alternative hypotheses about the intended interpretation of a natural language utterance in a given context, for use by the Dialogue Manager component of a discourse understanding system. In doing this SAD-93 performs, in addition to scope interpretation, interpretive processes such as reference resolution and tense interpretation.<sup>45</sup> Thus integration between scope interpretation and other aspects of pragmatic interpretation reflects the thesis defended throughout this dissertation that the scope assigned to operators is the result of the interaction of various discourse interpretation processes. SAD-93 is currently used as a module of the TRAINS-93 discourse understanding system.

<sup>43</sup> An example to the same effect is discussed in (Carter, 1987).

<sup>44</sup> The name stands for Scope And Disambiguation module, version 1993.

<sup>45</sup> SAD-93 only performs a limited amount of speech act interpretation, see below.

## 1.6 CONTENT OF THE DISSERTATION

In this dissertation I propose a theory of scopal ambiguity and an account of the process by which scopally ambiguous sentences like "I can't find a piece of paper" or "I don't see the butcher" are assigned a preferred interpretation in context. I argue against stipulating an independent 'scope disambiguation module'. I believe van Lehn was essentially correct when saying that 'people do not do scope'. I propose instead that both the preferences discussed in the literature on scope disambiguation, and those observed in the TRAINS corpus of natural language conversations, can be explained as the result of the interaction of discourse interpretation processes such as those devoted to the interpretation of referential expressions, of modals, or those concerned with relating the new utterance to the rest of the discourse. In the central chapters of this dissertation, Chapter 5 and Chapter 6, I present a detailed analysis of the discourse interpretation processes observed in the TRAINS conversations, and show how the scope preferences discussed above result from the interaction between these processes. The theory of discourse interpretation I adopt is not developed to account for the data about scope disambiguation, but is consistent with what is known about processes such as anaphora resolution, speech act interpretation, and the interpretation of definite descriptions.

I believe that the reason why this kind of approach hasn't been attempted before was the lack of tools in order to formalize the reasoning involved in discourse interpretation. In without assuming that the scope of operators has been determined, one needs to have a way to interpret syntactic structures. I propose a model of discourse interpretation as a process of inference over underspecified representations that is based on the DAS construction algorithm presented by Kamp and Reyle, but I provide an account of semantic ambiguity, and add the idea that perceived ambiguity is obtained when the inference process results in conflicting extensions.

The main technical tools I develop in the dissertation are (i) a new model-theoretic interpretation of semantic ambiguity, and a form of *underspecified representation* that captures this interpretation, thus underspecified differs in crucial respects both from the QLF forms used by Alhawi *et al.* and the Underspecified DRSs introduced by Reyle; (ii) the idea that dependency relations among information are best described in terms of relations between situations, and the relation of situation descriptions (the 'common ground' representation of part of the characterization of a situation) that I develop to capture this intuition; and (iii) the way I represent contextual dependency by means of *parameters*. This material is informally introduced in Chapter 3 and discussed in more detail in Chapter 4.

In Chapter 5 I review the syntactic and semantic treatment of the operators and constructs that appear in the TRAINS corpus and whose processes of interpretation I analyze in Chapter 6. In that chapter I propose a definition of the notion of 'operator', or scopally ambiguous sentence constituent, based on a distinction between sentence constituents that depend on the context for their interpretation and sentence constituents that don't. I also look at the impact of syntactic and semantic factors on the process of scope interpretation.

A system called SAD-93 has been developed that implements the theory of scope ambiguity described in this dissertation. SAD-93 is used as a component of the TRAINS-93 discourse understanding system developed at the University of Rochester. The implementation of SAD-93 is discussed in Chapter 7.

## 2 Formal Tools for Discourse Interpretation, Or: From Discourse Representation Theory and Situation Theory to Conversation Representation Theory

The theory of discourse interpretation proposed in this dissertation, *Conversation Representation Theory* (CRT), is a blend of ideas from Discourse Representation Theory, 'classic' Situation Theory as presented by [Barwise and Perry, 1983; Barwise, 1989; Devlin, 1991; Barwise and Cooper, 1993], and Episodic Logic, a 'conservative' version of Situation Theory developed by Len Schubert and Ching-Hee Hwang [Hwang, 1992; Hwang and Schubert, 1993]. The eclectic nature of CRT is motivated by the goals I intend to achieve.

- To provide a theory of semantic ambiguity and underspecification consistent with the Underspecification and Anti-Random hypotheses put forth in Chapter 1
- To substantiate the Condition on Scope Disambiguation presented in Chapter 1 by providing a simple theory of discourse interpretation as well as of discourse representation to be used to analyze the discourse interpretation processes observed in the TRAINS conversations and verify that indeed the observed scope preferences can be derived from the interaction of these processes. The development of such a theory requires a language in which those facts about conversations that play a role in discourse interpretation, such as their consisting of 'speech acts' uttered by different speakers, can be expressed

In this chapter I briefly review the three formalisms for semantic interpretation and discourse representation from which *Conversation Representation Theory* borrows the most. I only present the formal tools in this chapter; the actual analyses of semantic and/or discourse phenomena that come with these formalisms are discussed in the rest of the dissertation, whenever I propose alternatives. (This is especially true for Episodic Logic.) Those readers who are already familiar with these theories may wish to skip this material.



## 2.1 A CRASH COURSE ON DISCOURSE REPRESENTATION THEORY

Discourse interpretation depends on the state of the common ground. The common ground [Stalnaker, 1979; Heim, 1982] is "the participants' mutually developed picture view of what they are talking about" ([Chierchia and McConnell-Ginet, 1990], p. 166).<sup>1</sup> Discourse Representation Theory (DRT) is, first and foremost, a theory of the 'dynamics' of the common ground: how sentences modify the common ground, and how the common ground affects the interpretation of sentence. The DRT construction algorithm that specifies the impact of a sentence on the common ground is the best available example of a discourse interpretation procedure starting from an underspecified representation, and forms the basis for the discourse interpretation algorithm I propose in this dissertation.

It is not possible to provide more than a brief summary of DRT here. I refer the interested reader to [Kamp, 1981; Heim, 1982; Kamp and Reyle, 1993] for motivations and details. I use the version of DRT presented in [Kamp and Reyle, 1993].

### 2.1.1 'Donkey' Sentences and Discourse Representation Structures

Discourse Representation Theory was conceived to provide (i) a general account of the conditional (in) an account of the meaning of indefinite descriptions and (iii) an account of pronominal anaphora. The problems with the traditional theories of anaphoric reference and quantification addressed by DRT are demonstrated, first of all, by texts such as (2.1)

- (2.1) Pedro owns a donkey. He hates it.

According to the 'traditional' (e.g., Montaguean) treatment of indefinite NPs and reference, indefinites are existential quantifiers, and anaphoric reference to an indefinite is an example of bound anaphora. [Partee, 1972], a pronoun whose antecedent is an indefinite NP is logically equivalent to a variable bound by the quantifier. The text in (2.1) highlights a problem with this approach, namely, the fact that if we start by interpreting the first sentence, and we assign a scope to the existential quantifier, then we are unable to bind the pronoun in the second sentence.

A second difficulty with these traditional assumptions is illustrated by the so called 'donkey' sentences, i.e., sentences like (2.3) and (2.2). Again, we are interested in the interpretation of (2.2) in which the pronouns "he" and "it" are cases of bound anaphora, as opposed to referring to a contextually determined object. It is usually assumed that for an anaphoric expression to be bound by a quantifier, the anaphoric expression must be in the scope of that quantifier. Yet, there is no clear sense in which the indefinites "a man" and "a donkey" take scope over "he" and "it"—certainly the sentence doesn't have the interpretation that there is a man and there is a donkey such that if the man owns the donkey, then the man beats the donkey.

<sup>1</sup> According to Gross and Sidner [1996], the common ground is best seen as divided in three parts: information about the linguistic structure of the utterances in the dialog (the *Linguistic Structure*), information about the goals of the participants in the conversation (the *Intentional Structure*), and information about the objects introduced in the discourse and their relative salience (the *Attentional State*). I am concerned here with this part of the common ground that Gross and Sidner call attentional state.

- (2.2) If a man owns a donkey, he beats it.  
(2.3) Every man who owns a donkey beats it.

In fact, another problem with these sentences is to explain why indefinite NPs, that usually have an existential force (as in "A man owns a donkey") acquire universal force when embedded in a conditional.

The solution proposed in DRT (and arrived at independently by Kamp and Heim) was that indefinite NPs do not have a quantificational force of their own. Semantically, they behave as free variables, that can be bound by whatever operator they are in the scope of.<sup>2</sup> Let us consider first the case of an indefinite with existential force, like "a donkey" in (2.4). In the representation introduced by Kamp in [Kamp, 1981], the traditional formulas of first order logic are replaced by Discourse Representation Structures (DRSs). A DRS is a pair  $\langle M, C \rangle$ , where  $M$  is a set of discourse markers drawn from some set  $V$ , and  $C$  a set of conditions. For example, sentence (2.4) is represented in Kamp's version of DRT by the DRS in (2.5)

- (2.4) Pedro owns a donkey

$x, y$
PERSON( $x$ )
DONKEY( $y$ )
OWNS( $x, y$ )

- (2.5)

This DRS contains two discourse markers,  $x$  and  $y$ , and a set of atomic conditions: like DONKEY( $y$ )

Truth in DRT is defined in two stages. First, the notion of verification is introduced, that is analogous to the notion of satisfaction in ordinary logic. The verification conditions of the DRS in (2.5) are defined as follows: the DRS is verified by a variable assignment  $g$  with respect to a model  $M$  iff there is an assignment  $h$  that differs from  $g$  at most in the values assigned to  $x$  and  $y$ , and such that each condition in the DRS is verified by  $h$  (verification of atomic conditions is defined as satisfaction of conditions in standard first order logic).<sup>3</sup> A DRS  $K$  is then defined to be true wrt a model  $M = \langle U, F \rangle$  iff there is some assignment function  $f$  with values in  $U$  such that  $f$  verifies  $K$ . One can then define logical truth and entailment as usual.

From what I said above, it should be clear that the existential import of indefinite NPs is obtained in DRT by 'quantifying over assignments'. This ensures that indefinites are anaphorically accessible by all referential expressions evaluated with respect to the same variable assignment, such as pronouns in the same DRS or the representation of a text. The DRT representation for a text like (2.6) is obtained by just adding the conditions for the second sentence to the previously shown DRS representing (2.4), as shown in (2.7)

- (2.6) Pedro owns a donkey. He hates it.

<sup>2</sup> This scheme derives from the theory of unselective binding developed by Lewis in [Lewis, 1975].

<sup>3</sup> Kamp actually defines the semantics of DRS in terms of partial assignments and their extensions, but this is not important here.

$x$	$y$	$u$	$v$
PEDRO( $x$ )			
DONKEY( $y$ )			
OWNS( $x, y$ )			
$u = x$			
$v = y$			
HATES( $u, v$ )			

(2.7)

Quantifiers, including the 'implicit' adverbs of quantification in sentences such as "If a man owns a donkey, he beats it," are represented by means of complex conditions. A complex condition is a structure of the form  $\lambda_1 \rightarrow \lambda_2$ , where  $\lambda_1$  and  $\lambda_2$  are DRSs. For example, the translation of (2.2) is shown in (2.8).

(2.2) If a man owns a donkey, he beats it

$x$	$y$	$u$	$v$
MAN( $x$ )			
DONKEY( $y$ )			
OWNS( $x, y$ )			
$u = x$			
$v = y$			
HATES( $u, v$ )			

(2.8)

The satisfaction conditions for (2.8) are also defined in terms of quantification over variable assignments. A complex condition is satisfied by assignment  $g$  with model  $M$  iff for all variable assignments  $h$  that differ from  $g$  at most with respect to  $x$  and  $y$  and satisfy the DRS to the left of the arrow, there is a variable assignment  $f$  that differs from  $h$  at most in the values assigned to  $u$  and  $v$  that satisfies the DRS to the right of the arrow. The universal force of the indefinite is, again, achieved by quantifying over assignments.


## 2.1.2 The DRT Construction Algorithm

The DRT construction algorithm consists of a set of rules that map discourses into DRSs. The algorithm works as follows. It is assumed that a distinguished DRS exists, called the root DRS, this DRS is 'the common ground', as it were. The first step of the algorithm is to add the syntactic structure of the sentence to the root DRS. Syntactic structures are treated as 'special', uninterpreted conditions. Once the syntactic structure has been added to the root DRS, construction rules are repeatedly applied. A construction rule can be thought of as a production rule in rule based expert systems (e.g., Waterman and Hayes-Roth, 1978; Hayes-Roth, 1985; Walker 1987). It searches for a trigger—a certain syntactic configuration—and, if it can find such a trigger, it performs some operations on the common ground, such as adding new discourse markers and conditions, or replacing the syntactic structure that triggered the rule with a new syntactic structure. Most of the empirical input of DRT comes from the definition of the construction rules, and above all those for the interpretation of NP and of conditionals.

For an example of construction rules and of how the algorithm works, consider how it is used to obtain an interpretation for (1.54).

(1.54) Every kid climbed a tree.


The initial interpretation of (1.54) is the tree shown in (2.9).




The syntactic configuration in (2.9) triggers the construction rules for indefinites and universal quantification<sup>4</sup> shown below as CR INDEF and CR EVERY. I have adopted the rather self-explanatory rule presentation format adopted by Kamp and Reyle (Kamp and Reyle, 1993), with only slight modifications.

The rule CR INDEF should be interpreted as follows. The rule has first of all, a trigger. It applies whenever a DRS  $K$  (not necessarily the root DRS) includes an uninterpreted condition that contains an indefinite  $nt$  inside an 'S without any S's in between'. The specification of a construction rule includes the operations performed by the rule on the DRS to which it applies. CR INDEF has three effects: a new discourse marker  $x$  is added to the universe of  $K$ ; a new unary condition  $P(x)$  is added to the conditions of  $K$ ; and the syntactic configuration that triggered the rule is replaced in  $K$  with a new syntactic configuration, in which the NP has been replaced by the discourse marker  $x$ .

<sup>4</sup>Kamp and Reyle actually allow for a rule to have multiple triggers and specify all of the possible syntactic configurations. The notation with dotted lines allows for a simpler specification of the trigger, and is more general.





The rule for every-NPs, CR EVERY, is shown next. This rule is triggered by the occurrence in a DRS  $K$  of an uninterpreted condition in which an NP of the form "every  $P$ " occurs inside a  $S$ . The result is that the uninterpreted condition is replaced in  $K$  by a complex universal condition, that contains in its left-hand side (the restriction) a discourse marker  $x$  and a condition  $P(x)$ , and in its right-hand side (the nuclear scope) a new uninterpreted condition, obtained by replacing the NP "every  $P$ " with the discourse marker  $x$ .




Note, first of all, that this algorithm is based on the idea of underspecified representations, but the underspecified representations do not include semantic information and have no interpretation, therefore an interpretable DRS can only be obtained if all ambiguities are resolved.

Note also that the order of application of the construction rules determines the scope of the operations. If CR INDEF is applied first, the interpretation in which the indefinite takes wide scope is obtained; if CR EVERY is applied first, the other interpretation is obtained instead. Note that the purpose of the DRS construction algorithm is not, in general, to account for scoping preferences, but only to generate all the possible semantic interpretations, for this reason, the algorithm produces two partial interpretations for (2.9), shown in (2.10a) and (2.10b) (Versions of the algorithm incorporating syntactic-based scoping preferences can easily be implemented.)





The two 'partial hypotheses' in (2.10) are made into complete hypotheses by applying the remaining rules. Also, an underspecified condition that contains only a verb and discourse markers can be replaced by a predicate of the appropriate arity, whose arguments are the discourse markers. We thus obtain the two interpretations in (2.11)



The two other rules that have been used to build the DRSs discussed in this section are the rule for proper names and the rule for pronouns, that I briefly introduce below

**proper names rule:** if  $\alpha$  is a proper name, a new marker  $u$  ( $x$  in the example above) is added to the universal DRS (that is, the one not embedded in any other), and a new atomic condition of the form  $\alpha(u)$  ( $TREE(x)$  in the figure) is to the same universal DRS


**pronoun construction rule:** If  $\alpha$  is a pronoun, introduce a new marker  $v$  in the current DRS choose a 'uitable' marker  $u$  from the currently accessible ones, and add to the current DRS

a new condition  $v = u$  (A marker  $u$  is accessible from the DRS  $K$  if either  $u$  is local to  $K$ , or is introduced into a DRS which contains  $K$ , as discussed below)


### 2.1.3 Inference in DRT

Probably because of its unusual syntax and semantic properties, DRT is not used much either as a model of reasoning or in actual implementations of NLP systems.<sup>3</sup> Another problem has been solved in the past few years (Kamp and Reyle, 1991; Gabbay and Reyle, 1991; Sauer, 1991). I quickly review now the proposal of Kamp and Reyle in (Kamp and Reyle, 1991) to give an idea of how inference in DRT works

The system proposed by Kamp and Reyle is derived from Kalish and Montague's formulation of natural deduction. Suppose we have the DRS  $K$ :




and we want to prove the following extension of  $K$ ,  $K'$ :



inference is started by adding to (2.13) a new condition  $Show(K')$ , where  $K'$  is the part of (2.13) that is not included in (2.12)


<sup>3</sup> Examples of NLP systems that make use of DRT are the LILLOG system (Herzog and Rittinger, 1991) and the ACCORD system.

<sup>4</sup> This discussion is taken almost verbatim from Kamp and Reyle's own introduction



To prove the show line, inference rules are applied that introduce an alphabetic variant of the show line DRS. Once this is done, a Rule of Direct Proof licenses the elimination of the show line which means a successful completion of the proof.

The system of Kamp and Reyle includes three inference rules, among which the rule of Detachment (DET), that is a generalization of Modus Ponens. Detachment licenses adding to a DRS  $K$  a copy of the right hand side of a complex condition  $K' \rightarrow K''$  contained in  $K$  provided that the left hand side can be 'matched' with a part of  $K$ . In the example above the DRS  $\frac{x}{P(x)}$  can be matched by unifying  $x$  with  $z$ , and therefore the right hand side can be added. This transforms (2 14) in (2 15), which satisfies the Rule for Direct Proof and therefore represents a successful derivation.



## 2.2 SITUATION THEORY

Two major themes of this dissertation are that scopal preferences are the result of discourse interpretative processes that establish dependency relations (such as anaphoric relations) between 'information', and that these processes make use of pragmatic information about utterances as well as semantic information about sentences. Situation theory (Barwise and Perry, 1983, Landman, 1986b, Fensholt *et al.*, 1987, Barwise, 1989; Devlin, 1991) provides the conceptual tools for describing complex dependency relations between information that I use in the rest of the dissertation. Situation Semantics (Barwise and Perry, 1983, Gawron and Peters, 1990, Devlin, 1991) provides a framework for describing utterance interpretation. I quickly review the main ideas of Situation Theory and Situation Semantics in this section; the discussion is heavily modified in Devlin's book and in [Dekker and Hendriks, 1993].

### 2.2.1 Situations and Situation Inclusion

In all versions of Situation Theory, it is assumed that the 'real world' (or perhaps, our information about it) is 'carved up' in 'chunks' called situations:

Reality consists of situations—individuals having properties and standing in relations at various spatio-temporal locations (Barwise and Perry, 1983, p. 7)

The notion of situation in 'mainstream' Situation Theory is quite flexible and covers conceptual entities that in other approaches are distinct, such as events, states, etc. The event of Kim killing Lee is a situation, as are the state of Kim owning a car, a snapshot of what's going on in Rochester on August 3rd, 1993, etc. Situations are treated as primitives in the ontology with the same status as individuals, so the notion of real world situation is not really defined, but, loosely speaking, a real world situation is any aspect of the real world that we classify as a unit but we wouldn't think of as an object.

Situations in the real world are organized by a partial order relation of situation inclusion. the situation that consists of me going to the movies on February 19th, for example, includes sub-situations such as me driving to the theater at 7.45pm, me buying tickets at 7.55pm, me watching the movie from 8pm until 9.30pm, and me driving home afterwards. Each of these situations in turn includes sub-situations such as me asking the person at the window for a ticket, me handing the money to that person, etc.

### 2.2.2 Meaning and Information

Situations are carriers of information. Formal semantics is concerned with developing a theory of truth and truth-preserving inference. The goal of what I will refer to in this dissertation as 'mainstream' Situation Theory, that is, the work more directly inspired from [Barwise and Perry, 1983], is to provide an analysis of natural language based on information and the way information is extracted from the world and processed by human agents.

Situation Semantics begins with an analysis of the 'building blocks' of information—individuals and relations—and takes the notion of *information* ('information unit') as basic. Inform-

are primitive constituents of situations in this approach. An action is typically represented as a bracketed tuple consisting of a relation and a sequence of individuals:

((KULLIJDAM, sand))

In other approaches (e.g., Ilunduan, 1983b; Fenstad *et al.*, 1987b), *propositions* are used as the unit of information: a proposition is a (partial) function from situations to truth values that assigns 1 to a situation if a certain fact is a characteristic of that situation, 0 if that fact is not a characteristic of that situation, and is undefined otherwise. In this dissertation, I adopt propositions as the basic unit of information.

## 2.2.3 Situation Types

We can use information to classify situations. Each proposition effectively functions as a situation type: a real world situation  $x$  is of the type specified by a proposition if the proposition assigns to  $x$  the value 1. We can use this information to classify a certain situation as being of the same type as a previously encountered situation. For example, we may classify the situation occurring in the evening of February 19th as an instance of the situation type in which I go to the movies (a situation type with a large number of instances). This is usually written in Situation Theory using a notation like the following:

*(evening of February 19th) ⊆ GO-TO-MOVIES (maximal)*

(This expression reads: the proposition (inform) GO-TO-MOVIES (maximal) characterizes the real world situation evening of February 19th.) More complex situation types can be defined by combining propositions. For example, that particular evening may be classified as an instance of the type of situation in which I go to the movies and I see "His Girl Friday," also a situation type with a large number of instances, although not as many as the situation type of me going to the movies. Inform can be combined in the usual way by means of the standard connectives.

If we think of a situation type as specifying an "abstract" situation, we note that the formalism is actually rather flexible. Because any arbitrary collection of propositions may be interpreted as the specification of a situation type, the formalism has the capability for describing non-existing situations in addition to real-world ones. We can, for example, describe a possible (although not actual) situation in which I, on the evening of February 19th, I went to see a play.<sup>6</sup>

## 2.2.4 Partiality

Situation Theory also differs from traditional possible-world semantics in that a given situation, unlike possible worlds in the classical sense, can "keep silent" on the issue of whether

<sup>6</sup>In addition to the collection the action of proposition is used in mainstream Situation Theory as well—in fact two notations are used. Russellian propositions are the kind of proposition discussed above. Austinian propositions are statements of the form  $I_x \vdash \phi$  or equivalently  $I_x \vdash \phi$  that assert that situation  $x$  is of type  $\phi$ . I will consistently use the term proposition in order to Russellian propositions.

<sup>7</sup>A non-trivial, difficult question is what to do with non-world situation types.

a particular action/proposition is true. This can be expressed either by saying that the question of whether that inform is supported by that situation is not well-formed, or by saying that propositions are partial functions that may be undefined over certain situations.<sup>7</sup> There is then a sense in which situations behave like "small possible worlds" (Barwise, 1989).

This aspect of Situation Theory is crucial to the use of Situation Theory for assigning truth conditions to propositional attitudes, but leaves us with a number of decisions to make: for example, does a situation support the conjunction of two inferences if it supports each of the inferences?<sup>8</sup>

## 2.2.5 Situations Semantics

Situation Semantics is primarily concerned with the relation between natural language utterances, the context in which they occur, and the situations they describe. A number of concepts are introduced that turn out to be very useful for the purpose of describing the process of interpretation.

The basic idea of Situation Semantics is that whenever Kim utters the sentence<sup>9</sup>:

☞ Mary is running

two situations are involved. One is the situation  $w$  in which Kim makes the utterance. This situation is called the *utterance situation*.<sup>10</sup> The other situation is the situation  $d$  about which Kim is making a claim. This is called the *described situation*.

The utterance situation is often part of a *discourse situation* that consists of all the utterance situations that occurred during the conversation, and of an embedding situation—the part of the world in which the conversation is taking place.

The meaning of an assertive sentence according to Situation Semantics, is a relation between types of utterance situations, contexts and types of described situations:

$u, z \models \phi \parallel d$

Situation Theory started in part as an attempt to capture the truth conditions of perceptual reports, such as

(2.16) John saw Bill run

In *Scores and Other Situations*, Barwise argued that a simple way to capture the entailments of (2.16) is to make the object of seeing a situation. ☞ = notion of situation also provides a simple way to specify the semantics of tense: an intuitive formulation of the truth conditions

<sup>7</sup>I will follow the latter path below.

<sup>8</sup>The issue of partiality has raised a lot of interest. Two works that focus on the issue are: (Maddux, 1989) and (Langendoen, 1992).

<sup>9</sup>My presentation here is derived from chapters 4 and 8 of (David, 1991).

<sup>10</sup>I use in the dissertation the term *contextualized event* to refer to the same notion.

of a statement like (2.16) is that there is a situation in the past that supports the truth of the infinitive John see  $t$  where  $t$  is the situation of Bill running. In work starting with Barwise and Perry (1983) Situation Theory has also been concerned with modeling contextual dependency by means of parameters.

### 2.2.6 Parameters

The content determines those aspects of the interpretation of a sentence that vary from utterance to utterance. These aspects are represented in Situation Theory by means of parameters. A parameter  $x$  is an object that is part of the interpretation of a sentence, but needs to be anchored to some value in order for that sentence to carry information about a situation it represents an 'information hole', as it were.

Barwise and Perry proposed that pronouns and other referential expressions, such as definite descriptions, introduce parameters in the interpretation of a sentence. Their treatment of context-dependency was much extended by Ginzburg and Peters (1990). According to them, a use of the proper name 'John' introduces a parameter  $x_j$ —(NAMED  $x_j$  John)—where  $x_j$  is the resource situation that 'supplies a value for the parameter  $x_j$  the resource situation need not be the same as the situation at which the sentence of which the proper name is part is evaluated.

The notion of parameter plays a very important role in the formalization of discourse interpretation that I propose in the rest of the dissertation, although my interpretation of the notion differs somewhat from the one adopted by mainstream Situation Semantics.

## 2.3 EPISODIC LOGIC

Although Situation Theory has a number of appealing characteristics, it also represents quite a departure from common practice in formal semantics. Rather than adopting a wholesale, therefore, a number of researchers have been developing 'conservative' versions of Situation Theory. Such theories maintain the basic structure of Montaguean semantics, but incorporate ideas from Situation Theory, such as the notion of a structure of situations, or partiality. Theories in this vein are (Landman, 1990b; Fenstad *et al.*, 1987; Muskens, 1989; Kratzer, 1980).

Episodic Logic (Hwang, 1992; Hwang and Schubert, 1993) is one of these conservative theories. Episodic Logic was developed for semantic interpretation in natural language processing systems such as the TRAINS system that incorporates ideas from Situation Theory and Discourse Representation Theory.

### 2.3.1 The Language of Episodic Logic

The language of Episodic Logic is an extension of first order languages with restricted quantification. The basic language is fairly standard, if we ignore tense, the translation for (2.17), for example, is shown in (2.18).<sup>15</sup>

<sup>15</sup> Hwang and Schubert actually use an infinitesimal prefix for predicates. That is they write  $(\lambda x)(\text{period of time})x$ .

I adopt a more traditional prefix syntax here.

- (2.17) A dog came in.  
(2.18)  $\exists x \text{ DOG}(x) \text{ COMES-IN}(x)$

The language includes the usual connectives and determiners, including non-standard determiners such as 'the' or 'many'.

### 2.3.2 Truth as a Situation in Episodic Logic

The most relevant aspect of Episodic Logic is that terms can refer to situations. The language includes two operators to express truth at a situation, several predicates to describe temporal relations, predicates to describe situation inclusion, and a predicate  $\text{CAUSE}(r_1, r_2)$ : all of these constructs were primarily introduced to provide a treatment of tense and aspect.

Several of these constructs are displayed in the translation of 'John left' in (2.20).

- (2.19) John left  
(2.20)  $\exists x \text{ BEFORE}(x, \text{now}) \exists y (\text{LEAVE}(y, \text{John}) \wedge \text{e})$

(2.20) reads as follows: there is a situation (episode)  $e$  that takes place before the current moment in time now. The situation  $e$  can be characterized as John leaving. The language includes a set of temporal relations like  $\text{BEFORE}(x, y)$ , where  $x$  and  $y$  are situations or times.

The most important operator in (2.20) is the  $\text{e}$  operator, that represents truth at a situation. In order to introduce these operators, as well as the use of situations in Episodic Logic and how it relates to the notions introduced in §2.2, it is first necessary to say something about the semantics of Episodic Logic expressions.

The language of Episodic Logic is not typed.<sup>16</sup> Hwang and Schubert use a first-order model, in which the set of individuals  $I$  includes a set of situations  $S$  as a subset.  $S$  is a set of possible situations—it includes all the concrete situations obtained by taking 'chunks' of this world, as well as all the situations 'carved out' from other possible worlds. All of these situations are however pieces of an (actual or possible) reality. Each Episodic Logic formula denotes the characteristic function of a set of situations—what I have called above 'proposition' or 'situation type'.

The  $\text{e}$  ('single star') operator of Episodic Logic is used in expressions such as<sup>17</sup>

$$(\text{e} \star s)$$

to assert that the (possible or actual) situation denoted by  $s$  is an instance of the proposition denoted by  $\text{e}$ . The 'single star' operator may thus be interpreted as providing a *partial characterization* of a domain situation. The double star operator used in (2.20), instead, provides a *complete characterization* of a situation; that is, it asserts that the proposition denoted by the first argument of the expression  $(\text{e} \star \text{e} \star s)$  provides all the information about situation  $s$ . Double

<sup>16</sup> This aspect of Episodic Logic is not covered in the language I use, see below.

<sup>17</sup> A convention of Episodic Logic that I follow throughout the dissertation is to use square brackets to indicate expressions containing an implicit operator, such as  $\text{e}$ .





### 3 Conversation Representation Theory

**Conversation Representation Theory (CRT)** is the theory of discourse representation and interpretation that I have developed to serve as the foundation of my account of scope disambiguation. CRT is based on the three hypotheses presented in §1.5: the Underspecification Hypothesis, the Anti-Random Hypothesis, and the Condition on Scope Disambiguation. As discussed in §1.5, some consequences of these hypotheses are that (i) a distinction is made in CRT between semantic ambiguity and perceived ambiguity (in a particular approach to underspecification is adopted, illustrated by the logical forms used in §1.5), (ii) the assumption is made about discourse interpretation that processes such as reference resolution take place before the scope of operators has been determined, and (iv) perceived ambiguity is viewed as the situation which occurs whenever the result of discourse interpretation are distinct hypotheses.

This chapter and the next two are devoted to introducing Conversation Representation Theory in some detail, before presenting my hypotheses about the effect of discourse interpretation on scope (in Chapter 6). In this chapter I discuss the representation of the common ground in CRT and introduce the 'CRT interpretation algorithm'. In the next chapter, I propose a syntactic and semantic analysis of a small 'TRANS fragment'. The presentation in these two chapters will be mostly informal: the formal details of CRT are given in Chapter 4.

#### 3.1 FROM DRT TO CRT

In the Artificial Intelligence literature on discourse interpretation (e.g., Grosz, 1977; Wehner, 1979; Allen and Perrault, 1980; Grosz and Sidner, 1986; Carberry, 1990; Cohen and Levesque, 1990) it is assumed that prior to resolving references, identifying speech acts, and so forth the scope of operators has been determined. The formalisms used in that literature can then finesse the issue of how to represent information about surface structure of sentences, and concentrate on other kinds of information (e.g., about speech acts or about segmentation). Clearly, we can't make the same simplifying assumptions here. Situation Semantics gives a useful theory of what contextual information is available for interpretation, but not a story about how this information is brought to bear to obtain the interpretation of a new sentence.<sup>1</sup>

I argued in §1.5 that the hypotheses that discourse interpretation takes place before the scope of operators has been determined, and that in fact the results of discourse interpretation

<sup>1</sup>One notable exception is the work by Ginzburg and Peters in [Ginzburg and Peters, 1990].

are what determines the scope of these operators, suggest a theory of discourse interpretation that resembles DRT in many ways. What I mean is that a theory is needed that provides, first of all, a formal account of the content and organization of the common ground and its effect on the semantic interpretation of sentences. Moreover, the theory must be such that the input to discourse interpretation is an underspecified representation, and discourse interpretation is characterized in terms of 'rules' that operate on underspecified representations and, at the same time, update the common ground and result in more specific interpretations.

The DRS construction algorithm [Kamp and Reyle, 1993] is the only widely accepted proposal about the interpretation of discourses that, besides providing a detailed theory of what information is contained in a context, is not based on the assumption that the scope of operators has been determined. The language of DRT allows for 'unresolved conditions' that provide information about surface structure: the rules used by Kamp and Reyle are 'triggered' by the occurrence of certain syntactic patterns and result in modifications of surface structure forms.

As in DRT, it is assumed in CRT that the common ground contains information about (i) discourse markers, (ii) conditions on these discourse markers, and (iii) discourse marker accessibility. This information is represented by means of Discourse Representation Structures, or DRSs, that are (discourse marker, set condition set) pairs. It is also assumed that a distinguished DRS, called Root DRS, exists: this DRS represents all of the information about the common ground that is relevant for the present conversation.

The input to the CRT construction algorithm is an underspecified representation. The differences between the approach to underspecification taken in CRT and the one taken in DRT will be discussed below; what hasn't changed is the fact that discourse interpretation is cast in the form of a set of operations on DRSs like those discussed in §2.1, that may result in adding new discourse markers to a DRS, adding conditions, or deleting some. These operations can be 'triggered' by the occurrence of an underspecified representation as well.


The differences between CRT and DRT are motivated, first of all, by the goal of providing a theory both of semantic ambiguity and of perceived ambiguity, and also by the fact that, as a result of the Condition on Scope Disambiguation, CRT is meant to be not only a theory of discourse semantics, but a theory of discourse interpretation as well—and of interpretation as it occurs in conversations such as the TRANS dialogs, rather than in texts. The result is a theory of discourse concerned with both the *pragmatic* and the *semantic* aspects of the common ground.

### 3.2 THE INPUT TO DISCOURSE INTERPRETATION: LOGICAL FORMS

#### 3.2.1 Lexical Semantics and the Input to Discourse Interpretation

Although the CRT approach to discourse interpretation is inspired by the DRT construction algorithm, after seeing the examples discussed in Chapter 1 the reader will be aware of at least one difference between my approach to discourse interpretation and Kamp and Reyle's: the input to the CRT construction algorithm, unlike the input to the algorithm presented in [Kamp and Reyle, 1993], consists both of information about the s-structure of the sentence and of information about the semantic translation of the lexical items, this latter being the kind of

transition familiar from Montague Grammar.<sup>2</sup> For example, I mentioned in §1.5 that when processing the sentence "An engine must go to Avon," the first step in the CRT construction algorithm is to add the following logical form to the common ground



One of the implications of this proposal is that, while in DRT the semantics of lexical items is entirely specified by means of Model Construction Rules, there are two sets of rules in CRT, in addition to Model Construction Rules similar to those of DRT,<sup>3</sup> here also are semantic translation rules like those used in Montague Grammar or GPSG. The Model Construction Rules specify the 'dynamic' aspect of lexical items (i.e., whether they introduce discourse markers); the semantic translation rules specify the 'static' aspect of the semantics of the lexical items. But if adopting the form of underspecified representation used in Chapter 1 results in having two sets of semantic rules, one may wonder whether it would not be better to stick to the form of underspecified representation used by Kamp and Reyle.

There are good reasons for following the approach to underspecification adopted here, however. The first reason for believing that the input to discourse interpretation consists of more than just the information provided by s-structure (and the existing context) comes from the psycholinguistic literature: there is evidence that lexical interpretation takes place pretty early (Maratsen-Wilson, 1975; Hirst, 1987), and at least certain kinds of information that might be defined semantic, such as thematic role information, appear to affect parsing (Carlson and Tonenhaus, 1988). The same literature suggests that, although we have had to make some simplifying assumptions concerning the interface between parsing and the rest of discourse interpretation, there are good reasons not to do the further simplification of ignoring lexical information. Kutzman and MacDonald's data about the differences in scope preferences between perceptions predicates and agentive predicates, for example, indicate that information about the semantics of lexical items is used during scope assignment.

The second motivation for this approach to underspecification is that it is an unavoidable consequence of the Condition on Scope Disambiguation. If one accepts that hypothesis, then the input to discourse interpretation must include information as to which aspects of the semantics of the lexical items contained in an utterance must be resolved in context, until they aren't, the model construction rules cannot apply. And this information appears to be part of the semantic translation of the lexical items, rather than of s-structure. I have presented in Chapter 1 the

<sup>2</sup> To be more precise, the input to the DRS construction algorithm consists of a pair (not DRS s-structure) while the input to the CRT algorithm consists of a pair (not DRS underspecified representation). Below, whenever I talk of the input to discourse interpretation, I only consider the information coming from the last sentence, and ignore the rest of DRS.

<sup>3</sup> Except of course that their activation is controlled by the Condition on Scope Disambiguation.

arguments in favor of the Condition on Scope Disambiguation in task-oriented conversations; but even if one rejects the particular formulation that I gave of the hypothesis, it seems fairly clear that contextual resolution processes, as well, play a role in assigning a scope to operators. The contrast between sentence (3.1) (from Heim, 1982) and (3.2), for example, shows that scope assignment may be affected by reference resolution even when there is no salient topic in conversation.

- (3.1) If a dog barks at a cat, the cat (always) meows  
(3.2) If a dog barks, the cat (always) meows

The scope assigned to the definite NP "the cat" clearly depends on its discourse antecedent, thus in the preferred (for me) reading of (3.1), in which "the cat" is anaphoric on "a cat", "the cat" takes narrow scope with respect to the adverb of quantification "always", whereas in the preferred reading of (3.2), "the cat", interpreted as referring to a contextually salient cat, takes wide scope. This distinction cannot be captured by the DRT construction algorithm proposed by Kamp and Reyle, where the model construction rule for definites only makes use of syntactic information.<sup>4</sup> The algorithm could be modified by making it sensitive to coindexing at s-structure, and assuming that "a cat" and "the cat" somehow get coindexed at s-structure before the rest of the interpretation takes place—I am assuming that the input to the algorithm should be something like (3.3).

- (3.3)  $\lambda g \lambda g$  If a dog barks at a cat,  $\lambda g$  The cat, meows

While this simplification may be acceptable in many cases, we would still have to explain how it is that indices are introduced. And anyway, we can only get this far with syntactic coindexing; reference may in fact also be inferred by association, and it still affects scope.

- (3.4) When you buy a cat, make sure you check the tires

I don't know of any proposal to the effect that "the tires" should be coindexed with "the car" in (3.4), and it's hard to imagine how this proposal could be made to work. It seems fairly clear that to interpret "the tires" as a narrow-scope definite, we need to know that "tires" are objects found in "cars". In other words, to provide a general account of narrow-scope definites, one has to explain how definite descriptions are interpreted, and this process relies on more information than what one gets from s-structure.<sup>5</sup>

Yet another argument in favor of the hypothesis that more than s-structure is fed to discourse interpretation has been presented by David Dowty in [1986a]. Among the tasks of Kamp and Reyle's model construction rules is to establish temporal relations between the events in a discourse. As argued in [Hwang and Schubert, 1992; Kameyama et al., 1993], establishing

<sup>4</sup> Which interpretation is available depends on the version of the book—in the original drafts the rule CRD would only generate the wide-scope reading, as the final version, it only generates the narrow-scope reading. Kamp and Reyle are perfectly aware of the problem, of course—see the discussion at pages 297–300. I will also add that Heim's formulation of DRT in Heim, 1982 does not suffer from the problem.

<sup>5</sup> In fact, s-structure doesn't even provide the information that the interpretation of a lexical item is contextually dependent. Thus, if we take s-structure to be the only input to discourse interpretation, we are forced to say that interpretation refers to that which results in syntactic coindexing takes place after the model construction rules are applied.

these temporal relations requires information about the syntactic dominance relations existing in the sentence, thus this process has to operate before scopal relations have been determined. (See also Chapter 6.) However, it is also known (Parsée, 1984; Hinrichs, 1981), that the choice of temporal relations is affected by the aspectual class of the predicates involved (Dowty, 1979). For example, in (3.5a), a sentence whose main predicate is the telic TEL, is followed by a sentence whose main predicate is the telic LEAVE; the preferred interpretation is the one in which Mary leaves after John told her that Bill was a closet communist. In (3.5b), however, the first sentence is followed by a sentence whose main predicate is stative, LISTEN. Here the preferred interpretation is the one in which the situation described by the second sentence occurs at the same time as the situation of John telling Mary.

- (3.5) a John told Mary that Bill was a closet communist. She left.  
b John told Mary that Bill was a closet communist. She listened to him.

Accordingly, Kamp and Reyle make the model construction rules that establish these temporal relations sensitive to the value of a syntactic feature, [±STAT] ([1993], pages 541 ff.). Dowty, however, notes that, first of all, the choice of temporal relations is a default rather than a rigid norm, thus in sentences like (3.6) (from Lasnik and Asher, 1991), the event described by the second sentence temporally precedes the event described by the first sentence.<sup>4</sup>

- (3.6) Bill fell. John pushed him.

Secondly, Dowty notes that the aspectual class of a predicate can only be determined after semantic interpretation, and in fact, commonsense knowledge may often determine it, thus, while the predicate "build" is typically telic, the predicate "build houses" is atelic (see also Dowty, 1979; Mienis and Steedman, 1988). Both of these observations indicate that information about the semantic translation of lexical items must be available when determining the temporal relationship between events in a discourse.

I can't see of any way of addressing these observations that would result in a simpler system than the one presented here, and I will add one more observation: with the form of underspecified representations used in CRT, we can also assign an interpretation to partially underspecified representations, as shown below, whereas in standard DRT, only a completely disambiguated representation can be assigned an interpretation. The result is that a commitment to complete disambiguation is 'built in' the standard version of DRT, whereas in CRT, partially underspecified interpretations may be allowed.<sup>5</sup>

### 3.2.2 Adding Logical Forms to DRT

Augmenting DRT with logical forms like those seen in Chapter 1 is simply a matter of allowing DRSS to contain logical forms in addition to conditions like those used by Kamp and Reyle. Logical forms take the place of the uninterpreted conditions used in [Kamp and Reyle, 1993].

<sup>4</sup>For my knowledge, the best discussion of these data is in [Webster, 1987].

<sup>5</sup>A version of the DRT algorithm that also handles the aspect contains lexical information has been developed by Kamp and Raddecker (1992).

Several examples of logical forms have been shown in §1.5, but let's see one example in greater detail. Consider the sentence "Every kid climbed a tree," and ignore tense for the moment. In CRT, each lexical item—"every," "kid," "climbed," "a," and "tree"—has a semantic translation of the sort one finds in Montague Grammar.<sup>6</sup>

- "every"  $\sim \lambda P \lambda Q \lambda N x / s_1 \models P(x) / (Q(x))$   
"kid"  $\sim \text{KID}$   
"climbed"  $\sim \text{CLIMBED}$   
"a"  $\sim \lambda P \lambda Q / P(d) \wedge Q(d)$   
"tree"  $\sim \text{TREE}$

Let's assume for simplicity that the grammar consists of the phrase structure rules  $S \rightarrow \text{NP VP}$  and  $\text{VP} \rightarrow \text{V NP}$ .<sup>7</sup> Let's ignore the structure of noun phrases. We get the following logical form for the sentence "Every kid climbed a tree."

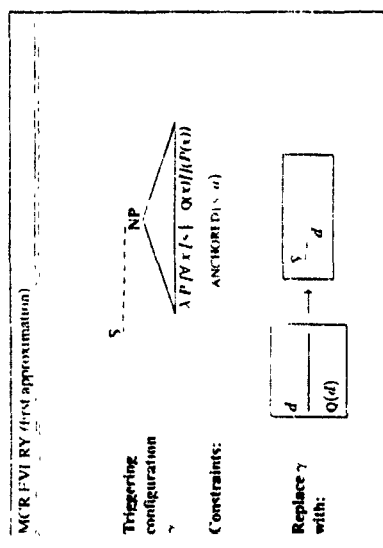
- (3.7)  $\text{IS} [\text{NP } \lambda Q \lambda N x / s_1 \models \text{KID}(x) / (Q(x))]$   
 $\text{VP} [\text{V CLIMBED}] [\text{NP } \lambda Q / (\text{TREE}(d) \wedge Q(d))]$

(3.7) is an s-structure whose leaves have been replaced by their semantic translation. Each internal node of a logical form such as (3.7) is labeled with a phrase category such as NP or VP, the leaves are labeled with expressions of CRT. Logical forms such as (3.7) can occur as conditions in DRSS; in first approximation, the first step of the CRT construction algorithm can be defined as the operation of creating a new DRSS out of an existing DRSS and a form such as (3.7).

In addition to the 'traditional' semantic translations above, there is in CRT a second set of rules, that I call *model construction rules*, whose role is to specify the 'dynamics' of discourse. These rules are operations on DRSS analogous to those one finds in DRT. The model construction rule for every-NPs, for example, is very similar to the rule CR EVERY discussed in §2.1. When applied, it replaces the triggering condition with a new tripartite condition, whose restriction contains a discourse marker that is only accessible from the nuclear scope, and whose nuclear scope contains a logical form obtained by replacing every-NP in the triggering condition with the new discourse marker. A simplified form of MC EVERY is shown below. (The rule will be discussed in more detail after having presented the semantic treatment of universal quantification.)

<sup>6</sup>There is nothing unusual in the semantic translation of the lexical items "kid", "climbed" and "tree", except that I distinguish lexical items from the predicates they translate using a small cap (not for the letter, as in CLIMBED, since of the 'youngest' notations favored by Montague). As discussed below, Milsark's distinction between *presuppositional* and *non-presuppositional* NPs plays an important role in CRT, and is captured by distinguishing between NPs whose translation includes a contextually dependent element represented by a parameter, such as "every kid", and NPs whose translation does not include a parameter, such as "a tree". The contextually dependent aspect of every-NPs is their *resource situation*—this is the way that domain restriction effects are captured in Situation Semantics [Barwise and Perry, 1983; Ginzburg and Perry, 1990; Devlin, 1991]. The parameter  $s_1$  in the translation of every stands for the resource situation. Individuals are translated as in DRT, as open formulas containing a discourse marker. Discourse markers are a special class of type  $e$  terms. The semantic translation of lexical items is discussed in greater detail in Chapter 5.

<sup>7</sup>In fact, in the syntactic framework I assume [Stowell, 1981; Pesetsky, 1982; Hagenman, 1991], the syntactic categories are slightly different, as I discuss in §3.3, but let's ignore this complication for the moment.



The point of this example is to show that the model construction rules of CRT differ from those of DRT because they are 'triggered' by the resolution of the contextually dependent aspects of the interpretation, represented in an underspecified representation by parameters such as  $x$ . The model construction rule for the universal quantifier applies if, and only if, the parameter  $x$  is 'anchored,' which means that its value gets fixed in context by discourse interpretation processes (see below).

In order to adopt logical forms as underspecified representations we need, first of all, a language that can be used to describe semantic translations of sentence constituents analogous to those proposed in the tradition derived from Montague, like those seen above. The approach followed here is to define CRT as a typed logic in the style of Intensional Logic, and to have in its language DRSs and the other constructs used in DRT. I discuss the semantic translation of lexical items, and the related model construction rules in detail in Chapter 5, the formal definition of CRT appears in Chapter 4.<sup>10</sup>

In fact, the real issue when dealing with underspecified representations is how to define their denotation, i.e. how to capture the notion of semantic ambiguity in a logic. This issue is discussed next.

### 3.2.3 Semantic Ambiguity and The Semantics of Logical Forms

In recent work by, among others, Alshawi and Crouch, Reyle, and myself [Presso, 1991, Alshawi and Crouch, 1992; Reyle, 1993], methods for assigning a denotation to expressions like (3.7) have been developed, which let us define logics for languages that include 'underspecified

<sup>10</sup> It is of course possible to develop a computational version of DRT. Urquhart and Klein (1986; Zeman, 1990) but the semantics translations assigned to lexical items in these systems do not differ significantly from those used in more traditional semantic systems.

expressions' in which a notion of *monotonic disambiguation* can be defined (the term is due to Alshawi and Crouch).

Both the system developed by Reyle and the system I propose in [Presso, 1991] assign to underspecified representations denotations that make them 'semantically equivalent' to a disjunction of the disambiguated readings. In my paper, for example, I adopt a 'relational' semantic system like those used by Heim in chapter 3 of her dissertation to account for the truth-conditional impact of anaphora [Heim, 1982] and by others [Barwise, 1987; Rooth, 1987; Schubert and Pelletier, 1988; Grendelink and Stokhof, 1991]. In these systems, the denotation of a logical expression is a relation, that is, a set of pairs of values, the elements of the pairs represent 'states': that is, sets of currently accessible anaphoric antecedents.<sup>11</sup> Since each of the independent interpretations of an ambiguous sentence denotes such a set, the denotation of an underspecified representation can be defined as the union of the denotations of its disambiguations. The result is that the sentence "Every kid climbed a tree" ends up denoting the same set of pairs as the sentence "It is either the case that a unique kid climbed all trees, or that for each tree, there was a kid that climbed that tree".

By defining the denotation of underspecified representations in this way, we get a logic of ambiguity that is very simple, and in which a notion of monotonic disambiguation can be defined. Yet, this definition doesn't truly capture the notion of semantic ambiguity introduced in Chapter 1. Take the sentence "He left", and assume that "he" may refer either to John or to Bill. If the goal of interpretation is to recover the *intended* meaning of a sentence, it seems clearly incorrect to attribute to the speaker of the sentence "He left" the intention of referring either to John or to Bill. Instead, what we would like to say is that it is either the case that the speaker intended to refer to John, or it is the case that she intended to refer to John.<sup>12</sup> I believe there are good reasons to distinguish such sentences from vague sentences such as "John is tall," where the speaker need not have intended to convey any information as to the precise height of John and may therefore be paraphrased to a certain extent as "John is taller than height X," which may be considered equivalent to a disjunction of assertions of the form "John has height Y".<sup>13</sup>

Alshawi and Crouch [1992] define the denotation of underspecified representations (that they call quasi-logical forms) using an approach that could be called 'supervaluations in reverse'.<sup>14</sup> They assume, first of all, that the denotation function  $\| \cdot \|$  is partial, that is,  $\| \phi \|$  may be equal to

<sup>11</sup> For example, the denotation of an existential formula  $\exists x P(x)$  is the set of pairs of assignments  $(f, g)$  such that  $g$  differs from  $f$  in that it assigns to  $x$  a value that satisfies  $P$ . I use the notation  $\| \phi \|$  in this way, if we have the test "A wig came in. It sat under the table," the value assigned to "a wig" is "accessible" when interpreting the premiss "it". On the other hand, states, operators such as the universal quantifier do not update the state, and therefore do not make new referents accessible. The denotation of the expression  $\forall x P(x)$ , for example, is a set of pairs of assignments of the form  $(f, f)$ .

<sup>12</sup> I shall briefly add that certain cases of pronominal reference with "it" do seem to be best treated as cases of vagueness rather than as cases of ambiguity. For example, we have in our dialogues numerous cases like (3.8) where "it" may refer either to the engine, or to the tractor, or both, but the intended meaning is not affected by the choice of the referent, given the effect of adverbs of focusing one object to the other.

(3.8) It-vent the engine to the tractor. Move it to Avon.

<sup>13</sup> In fact, it is dubious whether a 'disjunction' approach may be made to work even for vague sentences.

<sup>14</sup> A discussion of the supervaluation technique and how it can be used to provide a treatment of vagueness can be found in Chierchia and McCusker-Ginet's *Leitfaden* [1990] chapter 8.

0, 1, or be undefined. Second, they define  $\| \cdot \|$  in terms of an auxiliary relation  $W$  such that both  $W(\phi, 0)$  and  $W(\phi, 1)$  may be the case whenever that happens,  $\| \phi \|$  is undefined, otherwise, the value of  $\| \phi \|$  is either 0 or 1. The value assigned to a quasi-logical form  $\phi$  will be undefined if for one way  $\psi$  to disambiguate it  $W(\psi, 1)$ , whereas for another way to disambiguate it,  $\neg W(\psi, 0)$ .

The 'disjunction' fallacy is avoided in Alshawi and Crouch's system (an ambiguous statement is not necessarily equivalent to the disjunction of its disambiguated interpretations), but their approach still does not result in a clear distinction between semantic ambiguity and semantic vagueness. The same 'undefined' value may in fact be assigned both to a semantically ambiguous statement such as 'He left', in a context in which the pronoun 'he' may refer either to John or to Bill and only one of them left, and to a semantically vague statement such as 'John is tall', in case John is a borderline case of tallness. This is a symptom of a more general problem: by assigning the same 'undefined' value to all ambiguous statements, a considerable amount of semantic information is lost.

As it turns out, the definition of semantic ambiguity given in Chapter 1 leads to a formal characterization of semantic ambiguity that is not prone to the disjunction fallacy, yet does not result in loss of information: this is the idea that underspecified representations denote sets of propositions, namely, the set of all the propositions that correspond to semantically legitimate interpretations of a sentence. If we think of propositions as functions from situations to truth values, we conclude that logical forms should denote sets of functions from situations to truth values.

In fact, as we will see shortly, I propose that *all* meaningful expressions of type  $t$  should denote sets of functions from situations to truth values, and that the distinction between semantically ambiguous and semantically unambiguous sentences should be characterized in terms of the cardinality of the sets of functions that these sentences denote. Thus, the sentence 'He left', in a context where two possible antecedents for 'he' are available, will denote two functions from situations to truth values, the sentence 'Every kid climbed a tree' will similarly denote a set consisting of two functions—the proposition corresponding to the universal quantifier taking wide scope, and the proposition corresponding to the existential quantifier taking wide scope.<sup>15</sup> On the other hand, the sentence 'John is tall' will denote a single proposition, although the value of that proposition at some situations may be undefined if John is a borderline case of tallness.<sup>16</sup>

We may also arrive to the conclusion that underspecified representations should denote sets of propositions by a different route: namely, by asking ourselves what we can learn about the denotation of underspecified representations from methods for assigning a denotation to scopally ambiguous sentences such as the Cooper storage technique [Cooper, 1983]. The storage method

<sup>15</sup> As we will see below, Conversation Representation Theory is actually concerned with the interpretation of utterances, and the result of the 'two-step' translation is always an underspecified expression denoting a surface speech act. Thus the interpretation assigned to 'Every kid climbed a tree' actually consists of the proposition that is true at a discourse situation if the speaker of the sentence characterized the described situation as one in which the proposition resulting from the universal quantifier taking wide scope is true, and the proposition that is true at a discourse situation if the speaker of the sentence characterized the described situation as one in which the proposition resulting from the existential quantifier taking wide scope is true. The aspect of CRT will be discussed shortly.

<sup>16</sup> In fact, Alshawi and Crouch's denotation function can be defined in terms of the denotation function I am going to use—namely their function is one way of assigning a unique value to an expression  $\alpha$  in a situation  $\sigma$  if the set of values assigned by the propositions in the denotation of  $\alpha$  according to my notion of denotation

was developed by Robin Cooper as a way around a problem with Montague's quantifying-in technique, namely, the fact that in order to get all the readings of a scopally ambiguous sentence, one had to stipulate that the sentence was syntactically ambiguous (Thomason, 1974), see [Dowty *et al.*, 1981]. Briefly, Cooper proposes to define a function that assigns a value to syntactic trees, namely, to have their value be sets of sequences, each sequence representing a distinct 'order of application' of the operators that may result in an admissible interpretation of a sentence. For example, the quantifier 'every tree' can 'enter' the derivation of the VP 'climbed every tree' (whose logical form is shown in (3.9)) in two different ways: it is either possible to apply the translation of 'every tree' to the predicate immediately (which yields the narrow scope reading of the existential), or to apply the predicate to the variable quantified over and 'wait' before applying the quantifier (which is what happens then the wide scope reading is obtained).

$$(3.9) \quad \text{VP} [\text{y CLIMBED}] \\ \text{NP} \lambda Q \text{N} \lambda \text{f}_2 \models \text{TREE}(\text{N}) Q(\text{f}_2))$$

Cooper then proposes that the value of the NP 'every tree' be the set of two sequences shown in (3.10). One sequence consists of a single element, the 'traditional' Montague-style translation of 'every tree'. The second sequence consists of two elements: the result of applying the predicate CLIMBED to the variable  $x$ , and the semantic translation of the quantified NP, put 'in storage'. As a result, the value of the VP 'climbed every tree' consists of two sequences, as well: one obtained by applying the first element of the first sequence to the predicate CLIMB, the other obtained by applying the predicate CLIMB to the first element of the second sequence.

$$(3.10) \quad \{ (\lambda P \text{N} \lambda \text{f}_2 \models \text{TREE}(\text{N})) (P(\text{f}_2)),$$

$$(\lambda x \lambda P \text{N} \lambda \text{f}_2 \models \text{TREE}(\text{N}) / \text{CLIMBED}(x, \text{f}_2)) \} \\ (3.11) \quad \{ (\lambda x \text{N} \lambda \text{f}_2 \models \text{TREE}(\text{N}) / \text{CLIMBED}(x, \text{f}_2)), \\ (\text{CLIMBED}(\text{N}), \lambda Q \text{N} \lambda \text{f}_2 \models \text{TREE}(\text{N}) Q(\text{f}_2)) \}$$

The value of a sentence is then obtained as usual by combining the value of the VP with the value of the NP. The value of 'is Every kid climbed a tree' is also a set of two sequences, each representing a distinct reading of the sentence.

What does this tell us about the denotation of logical forms? There is an obvious way to assign a denotation to logical forms on the basis of the function CV that assigns to each logical form its 'Cooper value'. Loosely speaking, the denotation of a logical form  $\alpha$  can be defined as follows:

- Let  $\alpha$  be a logical form, and let  $\text{CV}(\alpha)$  be the set of single-element sequences  $\{(\sigma_1), (\sigma_2)\}$ . Then  $\| \alpha \| = \{ \sigma_1, \dots, \sigma_n \}$ .

In other words, that logical forms should denote sets of propositions is pretty much what one would expect, given the value assigned by CV to such forms.

I use Cooper's technique to define the semantics of logical forms in CRT. This has several advantages. First and foremost, Cooper discusses in detail how semantic and syntactic constraints on scope can be implemented by requiring that the storage be 'discharged' at certain positions—i.e., that the sequences 'carried above' those syntactic constructions that represent

barriers, such as  $\bar{S}$ , do not include any elements 'in storage'. This ensures that an quantifier in a clause may take scope over quantifiers in a higher clause, or in a sister clause, thus enforcing the Scope Constraint discussed in Chapter 1.<sup>17</sup>

We can also use the notion of storage to get a more precise definition of operator. In §1.1, I defined operators as those components of a sentence's meaning whose argument is not determined by  $\bar{s}$ -structure alone. This can now be rephrased more precisely as follows.

**Definition 3.1** An operator is a sentence constituent whose logical form representation (either a leaf or a maximal projection) has a Cooper Value consisting of two distinct sequences: one of which includes an element in storage, and the element in storage is the translation of either  $\bar{t}$  or  $\bar{H}$  or the Spec of the Logical Form.<sup>18</sup>

Note that operators are related to what in the government and binding literature would be called constituents subject to Quantifier Raising (May, 1985; Dieser, 1992).

I said above that all meaningful expressions of type  $t$  denote sets of propositions in CRT. This conclusion is actually forced upon us by the fact that ambiguity properties, a DRS that contains a logical form among its conditions is also semantically ambiguous, and therefore must denote a set of propositions as well. And since a DRS condition may have a DRS in its argument – in example are conditions that represent predicables that take semantic arguments, such as  $\text{is in } \text{Hill} \text{ or } \text{is in } \text{Hill}^2$  – these conditions must denote sets of propositions as well.

We are then left with an alternative: either having the language two kinds of conditions: one denoting sets of propositions, the other denoting propositions, or else letting all conditions (i.e., all expressions of type  $t$ ) uniformly denote sets of propositions. The second alternative besides being simpler, is also more general because it leaves open the possibility to deal with other kinds of semantic ambiguity—e.g., referential ambiguity and lexical ambiguity—and thus is therefore the solution adopted in CRT. All DRS conditions, and all other expressions of type  $t$ , denote sets of functions from situations to truth values, i.e., they have the same type of logical forms. The details of the semantics are discussed in Chapter 4. I will simply add here that the clause above defining the denotation of a logical form in terms of its 'Cooper Value' CV must be modified as follows:

- Let  $\sigma$  be a logical form, and let  $\text{CV}(\sigma)$  be the set of single-element sequences  $\{(\sigma_1, \dots, \sigma_n)\}$ ,  $\{(\sigma_1, \dots, \sigma_n)\}$ . Then  $\|\sigma\| = \{\sigma_1, \dots, \sigma_1, \dots, \sigma_n, \dots, \sigma_n\}$ .

<sup>17</sup> The definition of the function CV proposed by Cooper suffers from a technical problem: it generates spurious readings in the case of embedded NPs such as "a representative of every company", that have to be eliminated via a separated filter (Keller 1988). Keller proposed a technique that overcomes the problem (1988), more recently Perini (1999) argued that the right-scoping properties can be obtained without additional complications from the natural deduction approach to parsing. As I only consider here basic NPs that do not create problems for the simplest version of Cooper's technique, I use that here.

<sup>18</sup> The rather baroque way in which this definition is formulated has two reasons. First of all, an operator need not be a lexical item: for example, it is NPs that take scope, not just determiners. Secondly, we want to make sure that a (recursive VP whose object is a quantified NP) does not count as an operator.

<sup>19</sup> As discussed below, conditions of this form are useful to represent speech acts.

### 3.3 DESCRIBING PHRASE STRUCTURE

Up until now, I have been using a sort of vanilla description of phrase structure, based on a set of intuitive categories such as  $\bar{S}$ , NP, VP, etc. In fact, the theory of phrase structure I use in the rest of the dissertation is based on the version of  $\bar{X}$  theory developed in the 'principles and parameters' or 'GB' syntactic framework circa 1980 (Stowell, 1981; Chomsky, 1981; Chomsky, 1986a, 1986b; Haegeman, 1991). My main concern in this section is to highlight the differences between this approach to phrase structure characterization and the theories of phrase structure adopted in more traditional forms of context free grammar or in GPSG (Cardinaletti et al., 1985).

In the simplified form of  $\bar{X}$  theory developed by Stowell (1981), the grammar is reduced to three 'Phrase Structure Schemas':

- 1  $XP \rightarrow YP \bar{X}$
- 2  $\bar{X}' \rightarrow X \bar{Y}$
- 3  $XP \rightarrow X' YP$

The first two schemas say that a generic maximal projection NP is characterized by a maximal projection head  $\bar{X}$ , a maximal projection specifier YP (or [Spec:XP] for short) defined as the sister of  $\bar{X}$ , and a phrase structure and a maximal projection complement (also called [Comp XP]) defined as the sister of  $\bar{X}$  at phrase structure.  $\bar{X}$  can be any category. The third schema characterizes adjunction structures: a new structure of type XP can be obtained by adjoining a maximal projection YP to XP.

This theory of phrase structure is extremely simple, and therefore makes it very easy to define 'generalized' operations to combine semantic translations. It is also extremely restrictive: in particular, the following is assumed:



**Binary Branching Hypothesis:** All non-terminal nodes in a phrase structure tree have exactly two sisters.

Some obvious challenges to this hypothesis are coordination (see [Mollmann, 1992] for discussion and further references), double object constructions such as "John gave Bill a book" (see [Stowell, 1981; Larson, 1988] and there insertion sentences such as "There is a dog in the garden" [Roulland and ter Meulen, 1987]. The discussion of there-insertion sentences below should give an idea of the kind of analyses of these phenomena that have been proposed by Stowell and others to preserve the hypothesis.


Another characteristic of the syntactic framework adopted here is that sentences are assigned the structure in (3.12) instead of the most familiar structure in (3.13). The phrase category ' $\bar{S}$ ' is renamed as the maximal projection of a functional category T, for Tense.<sup>20</sup> Instead of a

<sup>20</sup> Functional categories are categories and lexically included in English, although they may be in other languages. The existence of categories such as Tense and Agr (for agreement) has been argued for since the very beginning of the transformational approach. It may be surprising, in particular, to learn that while in English tense is usually realized by means of inflections of the verb, in the framework adopted here it is assumed that tense occupies a position that dominates the verb as  $\bar{S}$  does. Tense is encoded with the verb as  $\bar{V}$ .

single *S*, a new maximal projection CP is introduced, whose head is empty in declarative matrix clauses, otherwise it is filled by a Complementizer—a lexical item such as “that”—in embedded sentences, such as “Kim believes that Sandy is stupid.”<sup>1</sup> The replacement of *S* with CP is motivated both by the desire to achieve a uniform phrase structure representation and by facts about the structure of questions



Recently the so-called ‘Split Infi’ hypothesis has often been adopted in the generative syntax literature (e.g. by Pollock (1989)). According to this hypothesis each functional category, such as Tense, Agreement or Negation, is the head of a separate maximal projection. This proposal is illustrated by the representation of ‘John does not like cars’ in (3.14).



This proposal is still under discussion. Besides using the category TP, I assume the existence of an (optional) maximal projection of negation NegP located above VP as in (3.14), but I omit AgrP. The structure of sentences is further discussed in §5.9.

### 3.4 CONTEXT-DEPENDENCY AND PARAMETERS

Because of the hypothesis that the scope assigned to operators is the result of processes that resolve the context-dependent aspects of a sentence’s interpretation, the context-dependency in underspecified representations is a central issue in this dissertation.

<sup>1</sup> The category complementizer was introduced by Bresnan in [Bresnan 1970].

In Situation Semantics (e.g., [Gawron and Peters, 1990]), the lexical items whose interpretation depends on context are classified as *parameters*: their interpretation depends on the value assigned in context to one or more parameters. Pronouns, for example, are translated as parameters, and the ‘domain restriction’ effect on quantifiers is captured by stipulating that the semantic translation of a quantifier includes a parameter to be contextually determined, the *reference situation*. Parameters in Situation Theory are special constituents of the universe.

As seen in Chapter 1 and in the previous sections of this chapter, parameters are used in CRT to translate contextually-dependent expressions of English as well; but although the ‘dotted’ notation of Situation Theory is preserved, semantically parameters are an entirely different type of objects than in Situation Theory.

The approach to contextual dependency adopted in this dissertation is based on the assumption that a ‘context dependent’ expression is semantically ambiguous, that is, it may have more than one interpretation, and more precisely, it is *referentially* ambiguous, which means that its value depends on the value of some other constituent of the common ground: the ambiguity comes from the fact that this constituent is not uniquely determined.

The denotation assigned to parameters in CRT is a consequence of this hypothesis, and of the general approach to semantic ambiguity taken in this dissertation. We have seen above that expressions of type *t*, that in a situation-theoretic reconstruction of Intensional Logic denote functions from situations to truth values, denote in CRT sets of such functions, some type *t* expressions of CRT denote singleton sets of propositions, while other expressions, such as logical forms and expressions that contain them, denote sets of greater cardinality. More generally, whenever a non-ambiguous expression of type *r* denotes an object  $\alpha$  in a logic like Intensional Logic, in CRT an expression of type *r* denotes a set of objects like  $\alpha$ , and semantically ambiguous expressions are defined as expressions that denote sets of cardinality greater than one. Other kinds of semantic ambiguity can then be treated in the same way as scopal ambiguity while in Episodic Logic an expression of type *e* denotes a function from situations to entities in the universe, in CRT these expressions denote sets of such functions, and referentially ambiguous lexical items can be made to denote sets of functions from situations to entities of the universe of cardinality greater than one. Parameters of type *e* are used to represent referentially ambiguous expressions, for example, i.e., a parameter of type *e* translating the pronoun “he”, denotes in CRT a set  $\{t_1, \dots, t_n\}$  of functions from situations to possible values of “he”.

The value of parameters depends on the current discourse situation: a parameter does not denote the set of all parameters from situations to objects in the domain, but the subset of those functions that map situations into objects that are ‘part’ of the current discourse situation, in a sense that will be discussed below and also in Chapter 4.

Referential ambiguity (context dependency) gets ‘resolved’ by *anchoring* a parameter. A parameter is *anchored* if only one among the functions in its denotation results in a consistent interpretation of the DRS in which the parameter occurs. A parameter can be anchored by means of equality statements of the form  $\{x = a\}$ , where *a* is not parametric, or it is already anchored: such equality statements make all but one of the interpretations of the parameter inadmissible. Once a parameter is anchored, it can be ‘replaced’ by a term that denotes the one function among those in the interpretation of the parameter that does not result in an inconsistent interpretation. In symbols, this can be expressed as follows. Let  $\text{PARAMETRIC}(x)$  be a logical predicate denoting a function that in a situation *s* maps objects of type *e* to 0 if the object denotes a singleton set,

to 1 if it denotes a set of cardinality greater than 1. We can then define the relation **ANCHORED** as follows:

$$\text{ANCHORED}(x, a) = \text{def } (x = a) \wedge \neg \text{PARAMETRIC}(a)$$

We can use the **ANCHORED** relation to formulate the following DRT style operation (schema) on DRSs:

PARAMETER-REPLACE	
Triggering configuration $\gamma$	$\Phi(x)$
Constraints:	$\text{ANCHORED}(x, a)$
Replace $\gamma$ with:	$\Phi(a)$

The goal of discourse interpretation is to anchor parameters. In fact, a conversation appears to be infelicitous unless all parameters can be anchored—the referents of all pronouns and definite descriptions have been identified the domain of quantification of all quantifiers has been appropriately restricted, and so forth—so much so that listeners appear to be ready to accommodate new information (e.g., introduce into the discourse some otherwise unspecified antecedent for a pronoun) rather than leave the interpretation parametric. This intuition, first noted by Lewis (1979) is formalized by the following condition:

**Condition on Discourse Interpretation** A discourse represented by the next DRS  $K$  is infelicitous unless all parameters occurring in  $K$  are anchored by the end of discourse interpretation

### 3.5 EVENTS, SITUATIONS AND DISCOURSE REPRESENTATION STRUCTURES

The concept of situation plays as important a role in CRT as it does in Situation Theory and in Fregean Logic. Just as these other theories, Conversation Representation Theory is based on the assumption that the 'world' can be partitioned into spatio-temporally coherent 'chunks'

called situations, and that a primary concern of natural language is to describe what holds at a situation.

This fundamental assumption motivates the fact, already discussed in the previous section, that all expressions of CRT denote (sets of) functions from situations to the value they would receive in *Intensional Logic*: sentence translations denote sets of functions from situations to truth values (propositions), expressions of type  $e$  denote sets of functions from situations to objects in the universe of discourse, and so forth. Also DRSs have a denotation of this type as I said in the previous section. CRT is a type-theoretic logic, whose language includes the constructs of DRT such as DRSs and complex conditions: a DRS in CRT is an expression of type  $t$  that denotes a set of propositions.<sup>27</sup>

#### 3.5.1 Situation Descriptions

As in Situation Theory and Episodic Logic, there is in CRT a construct to express the fact that a situation  $s$  is of the type specified by proposition  $\Phi$ , the expression  $s \models \Phi$ . These expressions are called *situation descriptions*. The semantics of situation descriptions is similar to that of the construct  $(\Phi \bullet s)$  in Fregean Logic and the construct  $(s \models \Phi)$  in Situation Theory, except that with situation descriptions the proposition that  $s$  is an instance of  $\Phi$  is specified by a DRS. The typical situation description is an expression of the form

$$(3.15) \quad \frac{\lambda x_1 \dots \lambda x_n}{s} \Phi$$

That discourse markers can be introduced as part of the description of a situation, rather than only in DRSs, is a very important aspect of CRT since it is used to model 'intrasentential accessibility', as we will see below. In addition to the expression  $s \models \Phi$ , where  $K$  may be a DRS, I also use at times the more syntax  $/s \models \Phi/$  from Situation Theory when the proposition  $\Phi$  is not a DRS.

I make a rather extensive use of the indexical  $s^*$  in the dissertation. This indexical term refers to the situation of evaluation, i.e., the value of  $s^*$  at a situation  $s$  is  $s$  iff

#### 3.5.2 Events and Event Descriptions

In addition to situation descriptions, there are in CRT expressions of the form  $e \models \Phi$ , called *event descriptions*. The distinction between events and situations does not reflect a difference in ontological status, but an informational difference. I use the term 'event' to refer to a situation that is introduced into discourse as the translation of a single finite clause, and is completely characterized by the information provided by the translation of that clause.<sup>28</sup> For example, the

<sup>27</sup> Loosely speaking, the 'situation' that assigns 1 to a situation if all conditions are true at that situation for some assignment of values to discourse markers. I return on this issue in more detail in Chapter 4.

<sup>28</sup> Perhaps a bit confused are the term event for situation characterized by both teleic and atelic predicates.



sentence "Kim greeted Sandy" introduces into the common ground an event—a situation that can be completely described by the fact that Kim greets Sandy. I use 'situation' both as a move general term and to refer to situations about which the common ground contains only partial information.)

The difference between situation descriptions and event descriptions is similar to (and in fact, derives from) the distinction between the expressions  $\{\Phi \leftrightarrow s\}$  and  $\{\Phi \leftrightarrow s\}$  in  $\mathcal{F}$  predic Logic: situation descriptions are *partial* characterizations of a situation, whereas event descriptions provide a *complete* characterization. In other words, the following schema, where  $\Phi$  and  $\Psi$  are propositions

$$\exists s (s \leftrightarrow \Phi \wedge s \leftrightarrow \Psi \equiv \Phi \equiv \Psi)$$

is consistent in CxT, but the following schema is a logical contradiction

$$\exists s (s \leftrightarrow \Phi \wedge s \leftrightarrow \neg \Psi \equiv \Phi \equiv \neg \Psi)$$


Another difference between situation descriptions and event descriptions is that only situation descriptions allow discourse markers to be accessible outside the DRS in which they are introduced, as we will see below.

### 3.5.3 Situation Kinds

As it turns out, participants in conversations do not refer only to actual situations: in the TRANS conversations, for example, it is common for the user and the system to refer to objects that are part of the plan and/or to events that are included of the plan, although the plan is not a part of the real world. As discussed in §2.2, Situation Theory allows for *abstract* situations in addition to actual situations. Abstract situations are not really situations, but collections of intensions, that can be used to classify actual situations: an actual situation is an instance of an abstract situation if it supports all the intensions in the abstract situation. In Episodic Logic, functions from situations to truth values, or 'situation types', can similarly be used to classify situations. In addition, Episodic Logic allows for the *nominalization* of situation types (and other predicates) by means of a *kind-forming* operator  $K$ .<sup>24</sup>

There are reasons to believe that situation kinds are the appropriate way to model the way plans are referred to in conversations (for some of these arguments, see §5.8 and §6.1). Since the situation types of Episodic Logic correspond to what I have been calling propositions, we can get situation kinds by stipulating that the model contains a set of kinds, such that an isomorphism  $\kappa$  from propositions to situation kinds can be defined. We can then define  $K$  as the object language realization of  $\kappa$ . Situation kinds are introduced in the common ground by equality conditions of the form

<sup>24</sup>The kind-forming operator is closely related to analogous operators proposed as in property theory by Chierchia and Turner (Chierchia, 1984; Chierchia and Turner, 1988). Chierchia (1984, p. 18) attributes the proposal that Carlson's kinds are predicates nominalizations to Coccia and Pietri (1976); a similar proposal was also made in [Schubert and Pelletier, 1988].



I also use the *kind-instantiation* relation introduced by Carlson [1978]:  $i \text{ is } R \text{ of } k$  if the object  $i$  is an instance of kind  $k$ . Finally, a SUBKIND relation is needed that mirrors the subset relation between properties: in the sense that kind  $k$  will be a subkind of kind  $k'$  (the value of  $\text{SUBKIND}(k', k)$ ) iff for each object  $a$  in the denotation of  $\kappa^{-1}(\{k'\})$ ,  $a$  is also in the denotation of  $\kappa^{-1}(\{k\})$ . This can be expressed in the object language using  $K$  as follows

$$(3.16) \quad \forall k_1, k_2, \text{SUBKIND}(k_1, k_2) \equiv \forall s [s R k_2] \supset [s R k_1]$$

## 3.6 PRESUPPOSITION AND THE INTERPRETATION OF NOUN PHRASES

As said before, a key idea in CRT is that the goal of discourse interpretation is to 'fill in' those aspects of interpretation that are left underspecified by syntactic and lexical interpretation; these components of the meaning of a sentence translate as parameters. The predictions of the theory therefore crucially depend on which sentence constituents are assumed to be parametric, and which aspect of their interpretation is taken to be context-dependent.

In this section I argue that the classification of noun phrases into *strong* and *weak*, discussed by Miltsark [1977], Barwise and Cooper [1981], and in the papers in [Reuland and ter Meulen, 1987], tells us which NPs have a context-dependent meaning. Following de Jong [1987], I propose that strong NPs are those that presuppose the existence of a set of objects of the type specified by the restriction that is distinct from the set of all individuals in the domain. (An alternative way to say this is to say that strong NPs are those whose *domain of quantification* has to be identified.) As in Heim [1982; 1983b], presupposition interpretation (in particular, the dual processes of *accommodation* and *cancellation*) are formulated in terms of context. My approach to existential presupposition is close to that proposed by van der Sandt, who sees presupposition as anaphora, but I reformulate his proposal in terms of parametric objects.

### 3.6.1 Strong and Weak NPs

In the literature on quantification—e.g., [Miltsark, 1977; Barwise and Cooper, 1981; Diesing, 1992]—a distinction is traditionally made between *strong* and *weak* NPs. Miltsark's original distinction was motivated by the distribution of NPs in post-verbal position in there-insertion contexts. Weak NPs, such as "a boxcar", "some boxcars", or "many boxcars", are allowed to

appear in these contexts, whereas strong NPs, such as "every boxcar", "most boxcars", "the boxcar", and "it", cannot occur in that position. The contrast is exemplified by (3.17)

- (3.17) a \*?There is (every, the, that) boxcar in the station  
b There are (some, two, many) boxcars in the station

Miltsark proposed that strong NPs only have a quantificational interpretation, whereas weak NPs are ambiguous between a quantificational and a cardinal reading, only cardinal NPs are allowed in there-inversion contexts

Barwise and Cooper [Barwise and Cooper, 1981] formulated the distinction between strong and weak NPs in model-theoretic terms. They define a determiner Det as (i) *positive strong*, (ii) *negative strong*, or (iii) *weak* according to whether the sentence "Det N's are N's" is tautological, contradictory, or contingent, respectively. Thus, "every" is positive strong, as "Every boxcar is a boxcar" is tautological, "no" is negative strong, as "No boxcar is a boxcar" is contradictory, and "three" is weak, as the truth of "Three boxcars are boxcars" depends on there being three boxcars in the domain.<sup>2</sup>

De Jong [de Jong, 1987] reconstructs Barwise and Cooper's distinction between strong and weak NPs as a distinction between determiners that presuppose that the denotation of the property denoted by their restriction is not empty, and determiners that do not carry such presupposition. De Jong argues that the true logical behavior of strong determiners is displayed in sentences such as "Every kid entered the room," rather than in sentences such as "Every kid learns to walk before he learns to talk", and that in the latter case, the lack of presupposition is due to the generativity of the sentence.

I follow De Jong's proposal, that is closely related to proposals about representing domain restriction in Situation Theory (see, e.g., Gawron and Peters, 1990, Croper, 1993). Before introducing my proposal, I will discuss my treatment of existential presuppositions.

### 3.6.2 Presuppositions

Presuppositions are an important part of the meaning of lexical items. Proposition  $\Phi$  presupposes proposition  $\Psi$  iff the truth of  $\Psi$  is vital to establish whether  $\Phi$  is the case. I.e., if  $\Psi$  is not the case, one cannot rightly ask whether  $\Phi$  is or is not the case. Presuppositions are most commonly tested by checking whether they follow both from a statement  $\Phi$  and from its negation  $\neg\Phi$ .

Well-known examples of lexical items whose meaning is in part defined in terms of presuppositions are definite descriptions and verbs such as "regret". Strawson [1950] claimed that sentence (3.18) does not appear to be either true or false in case there is no king of France, *contra* Russell [1905]. While this claim has been repeatedly challenged (e.g., by Neale [1990]), there is little dispute that presuppositions are involved in the meaning of "regret". (3.19a) presupposes, rather than entailing, that John told Mary about Bill's misadventure, as shown by the fact that the same inference can be drawn by its negation (3.19b).

- (3.18) The king of France is bald

<sup>2</sup> An alternative model-theoretic formulation is proposed by Keenan [1987]

- (3.19) a John regrets telling Mary about Bill's misadventure  
b John doesn't regret telling Mary about Bill's misadventure

Presupposition accommodation and its converse, usually called *cancellation* [Karttunen and Peters, 1979, Gazdar, 1979] are the most important aspect of presuppositions, as far as discourse interpretation is concerned. The problem of accommodation is the problem of adding to a context what is presupposed by a sentence in the context. The problem of cancellation is formulated as follows: while (3.19a) presupposes John telling Mary about Bill's misadventure, (3.20) doesn't. The presupposition gets 'cancelled' somehow.

- (3.20) If John told Mary about Bill's misadventure, he regrets telling Mary about Bill's misadventure

The differences between the accounts of presupposition that have been presented in the literature usually lie in how presuppositions get cancelled. What strikes me as the most plausible among these accounts is Heim's theory, that I discuss in the next section.

### 3.6.3 Presupposition as Anaphora

Heim proposes a context based account of presupposition. Her definition (Heim, 1983b, p. 117) is as follows

S presupposes p iff all contexts that admit S entail p

Heim identifies 'context' with 'files', a notion introduced in chapter 3 of her dissertation that is closely related to the notion of DRSs (a file denotes a set of *(assignment world)* pairs, much as a DRS world), and uses this definition to explain presupposition cancellation in examples like (3.20) as follows: REGRET(<sub>i</sub>, TELL-ABOUT(<sub>j</sub>, *my bill-misadventure*)) (let's call this  $\Phi$ ) presupposes TELL-ABOUT(<sub>j</sub>, *my bill-misadventure*) (let's call this  $\Psi$ ) iff every file (assignment) that admits the former entails the latter. The semantics of the conditional (3.20) in File Change Semantics and DRT is such that in an assignment that satisfies the conditional, every assignment that satisfies the antecedent can be extended to an assignment that satisfies the consequent. The semantics of the conditional thus guarantees that every context of evaluation that satisfies  $\Phi$  (locally) entails  $\Psi$ , and therefore, there is no 'global' requirement on the context of evaluation to entail  $\Psi$ .

According to van der Sandt, as well [van der Sandt, 1988; van der Sandt, 1990] there is an such a thing as 'presupposition cancellation'. He proposes an analogy with anaphora. He notes a parallelism between environments that allow for anaphora, like those in (3.21), and environments that 'cancel' presuppositions, like those in (3.22). Van der Sandt notes that the parallelism holds for all types of presuppositions—definite descriptions, factives, clefts, presuppositional adverbs, and so forth.


- (3.21) a John owns a donkey. He beats it.  
b If John owns a donkey, he beats it.  
(3.22) a Jack has children. All of Jack's children are bald.  
b If Jack has children, all of Jack's children are bald.

Van der Sandt proposes then that the presuppositions of a sentence are always added to the common ground, when the content of the presupposition is already part of the common ground (a suitable antecedent is found<sup>26</sup>), the two facts are identified: otherwise, the context of the presupposition is accommodated. The difference between anaphora and presupposition, according to the van der Sandt, is that the latter includes enough material that accommodation is always possible.

Van der Sandt also proposes in [van der Sandt, 1990] an algorithm for presupposition accommodation and cancellation. This algorithm crucially relies on treating DRSs as representations. In the next section I propose a formulation of the notion of existential presupposition in terms of parameters and situation descriptions that, besides making the parallel between anaphors and presuppositions even more evident, does not require assuming that DRSs are syntactic structures.

### 3.6.4 Presupposition, Parameters and Situation Descriptions

The separation between assertional content and presuppositional content, as well as the results obtained by the algorithm for presupposition accommodation proposed by Van der Sandt<sup>27</sup> can be achieved by assuming that the existence presuppositions associated with sentence  $S$  are statements about a situation that is not included in the situation described by  $S$ , and whose value has to be provided by context. In other words, I propose that the semantic translation of definite descriptions, and other lexical items whose meaning includes a presuppositional aspect, includes a parametric component. The truth conditions of the statement "The King of France is bald" are as shown in (3.23):<sup>28</sup>




This representation includes a part that must be true at the situation of evaluation, namely, the fact that the object denoted by  $x$  in the situation denoted by  $s$  must be bald, and an existence presupposition represented by the embedded situation description. The evaluation of  $x$ , the situation described by  $S$ , depends on there being a situation  $s'$  that includes a King of France, that can be used as an anchor for the parameter  $x$ . The existence of a king of France is not part of the statements about  $s$ . On the other hand, a context of evaluation (the DRS in which the statement is embedded) must provide a value for  $x$  in order for the statement to be interpretable, just as in the case of any other parameter, as requested by the Condition on Discourse Interpretation in §3.4.

The negation of the sentence gets the same presupposition


(3.24) The King of France is not bald

<sup>26</sup> That for example allows for the possibility of ambiguity in presupposition cancellation.

<sup>27</sup> I discuss the translation of definite descriptions in more detail in §5.1.




This proposal is closely related to Heim's and van der Sandt's. If the context (in our case, the local DRS, or one of the DRSs accessible from it) does not include a statement to the effect that situation  $s$  supports the fact that there is an  $x$  such that  $x$  is the King of France, the fact has to be accommodated, as proposed by Heim and van der Sandt. This is impossible in case the context already contains a statement to the contrary, i.e., that there is no King of France: that is, if the context contains a statement of the form:



On the other hand, if the context already contains such a statement, the parameter gets 'anchored' in the situation that supports it, thus the presupposition gets canceled much as proposed by Heim and van der Sandt. Consider the conditional in (3.27) (from Heim's thesis). The representation I propose for the conditional is shown in (3.28).

(3.27) If a dog barks at a cat, the cat meows.



When the NP "the cat" is interpreted as anaphoric on  $y$ , the parameter  $s$  is anchored to  $s'$ ; in other words, there is no need to accommodate a situation including a cat in order to interpret the statement.<sup>29</sup> When no such anaphoric relation can be established, the presupposition can be accommodated by introducing a new sub-situation of the discourse situation  $s'$  and letting  $s$  be identical with  $s''$ .

<sup>29</sup> The interpretation of definite descriptions is discussed more in detail in §6.5.

### 3.6.5 Existential Presuppositions and Strong Noun Phrases

I propose that strong noun phrases are defined as those noun phrases whose semantic translation is parametric, that is, depends on context. Weak NPs, on the other hand, are those noun phrases whose translation does not depend on context. Thus, the indefinite "a boxer" receives an Heimian translation as follows.

$$^a \text{boxcar} \sim \lambda P \text{BOXCAR}(d) \wedge P(d)$$

where  $d$  is a discourse marker, discourse markers, as explained in §4.3, are CMT expressions that behave like unbound variables in DRT. With Milnerk (1974), Diesing (1992) and others, I assume that certain determiners, like "some" and "many", are ambiguous between a weak and a strong interpretation, under the strong interpretation, they denote a relational determiner treated like the other strong determiners below.

Roughly speaking, I assume that there are three classes of strong NPs, depending on which aspect of their translation is parametric

- *primaries* are essentially parametric, i.e. "the translations  $\lambda x.b$
- *proper names* and *deictic descriptions* introduce a discourse marker into context (as in DRT) but they also introduce parameters. In the case of definite descriptions, two parameters are involved: the resource situation and the object in the discourse that the definite description is associated to

$$^a \text{the boxer} \sim \lambda P / s \models \text{BOXCAR}(d) \wedge [d = x]$$

This translation is closely related to Heim's translation for definites, as in Heim's proposal, the translation for definites differs from that for indefinites in that (i) the content of the restriction is presuppositional, and (ii) the discourse marker has to have the same value assigned to another discourse marker in context.<sup>20</sup>

- In the case of *strong quantifiers*, the parameter is the resource situation. Thus, "every boxer" translates as follows

$$^a \text{every boxer} \sim \lambda P / s \models \text{BOXCAR}(x) \wedge [P(x)]$$

What this translation says is that in the case of this NP,  $x$  has to be identified in the context in a situation such that all boxcars in that situation have the property  $P$ . This, of course, corresponds to performing a domain restriction on the quantifier.

<sup>20</sup> In work on progress, I argue for a more general treatment of definites in that  $\text{BOXCAR}(x)$  defines like "the boxer" is a context in which a boxer has been introduced and he is interpreted. The translation proposed here differs from § 3.6.1 in that it says in that  $x$  you find that a relation must exist between the new discourse marker and a discourse marker in context. This relation must be equalities as a special case.

I discuss all of these translations in more detail in Chapter 5. It is a good idea, however, to defend right away the treatment just sketched against two possible attacks. First of all, it may look like according to this treatment, NPs like "every boxer" are predicted to be as presuppositional as, say, NPs such as "the boxer" or "both boxcars". That's not the case: these latter NPs require that *both* the resource situation and the 'referent' of the NP be identified, while NPs such as "every boxer" only require the resource situation to be identified. Thus, accommodating "every boxer", when the context contains no possible antecedents, will be much easier than accommodating "the boxer".<sup>21</sup>

Secondly, one might argue that sentences such as "Every chess player hates donkeys" do not presuppose the existence of chess players. However, it has been argued by Cooper in [1993] (p. 206-207) and by Kratzer in unpublished material quoted by Diesing [1992], that this is just another case of presupposition cancellation: the resource situation of the quantifier gets bound by the generic operator that is part of the meaning of generic sentences. Thus, "every chess player hates donkeys" would be interpreted as "in every normal situation, every chess player in that situation hates donkeys".<sup>22</sup>

<sup>21</sup> Because of the possibility of accommodation, the hypothesis about a sentence obtained at the end of interpretation will be equivalent to the interpretation proposed by Heim and Peters [1990] and Cooper [1993], who propose that the resource situation is always a discourse marker—like, for example, can be unselectively bound by generic operators. This alternative, however, leaves us without a way to characterize the difference between strong and weak NPs.

<sup>22</sup> Eng also arrived at the conclusion that strong NPs are all and only the NPs with existential presuppositions [Eng 1991] although her treatment of existential presuppositions differs from mine. One of the arguments she brings in support of her theory is the fact that in Turkish, where specific NPs are morphologically marked (they receive Acc case), all strong NPs are marked as strong (p. 10 and 11).

### 3.7 DISCOURSE INTERPRETATION AND THE PRAGMATICS OF DISCOURSE

In this section I briefly discuss an important discourse interpretation process occurring in the TRANS conversations, the interpretation of definite descriptions. I do not present any new data on the facts about definite description interpretation. I discuss here all have been previously studied in work on discourse interpretation in Artificial Intelligence and Linguistics, such as [Wether, 1979; Grosz, 1977; Cohen, 1978; Carberry, 1980; Grosz and Sidner, 1986]. The goal of the discussion is to show the impact on discourse interpretation of pragmatic factors such as the organization of utterances in discourse segments and the position of the current focus of attention. The conclusion I intend to draw is that a new view of the common ground as a model of the discourse situation in which the conversation takes place, and in which pragmatic information is represented, is necessary to develop an account of discourse interpretation and to model conversations.

I then propose that the impact of pragmatic factors on interpretation can be accounted for by reinterpreting the role of the root DRS and discuss how certain forms of information about the conversation can be represented in a root DRS with this kind of interpretation.

#### 3.7.1 Definite Description Interpretation and the Common Ground

The main facts about the use of definite descriptions in the TRANS dialogues can be illustrated with reference to the (edited) fragment in (1.61) repeated here for convenience

- (1.61)
- 13.1 U not at the same time
  - 13.2 okay
  - 13.3 We're gonna hook up engine E2 to the boxcar at Elmira
  - 13.4 and send that off to Corning
  - 13.5 now while we're loading that boxcar with oranges at Corning.
  - 13.6 we're gonna take the engine E3
  - 13.7 and send it over to Corning.
  - 13.8 hook it up to the tanker car.
  - 13.9 and send it back to Elmira
  - 14.1 S okay
  - 14.2 We could use one engine to take both the tanker and the boxcar to Elmira
  - 15.1 U oh we can do that?
  - 16.1 S yeah
  - 17.1 U then bag the whole thing with engine E3
  - 17.2 and just hook up the tanker car with the boxcar that has oranges in it
  - 17.3 and take it back to Elmira
  - 18.1 S okay
  - 18.2 that's no problem
  - 20.1 U okay
  - 20.2 great

- 29.3 while this is happening.
- 29.4 take engine E1 to Danville.
- 29.5 pick up the boxcar.
- 29.6 and come back to Avon
- 30.1 S okay
- 31.1 U okay
- 31.2 there load the boxcar with bananas

Of the eight major usage types of the definite article in English [Hawkins, 1978], two are especially common in our transcripts. The two uses of the definite description "the boxcar" in (1.61), in 13.3 and 29.5, are instances of *visible situation* use of definite NPs, which occurs when "...the object referred to is visible to both speaker and hearer in the situation of utterance, and is furthermore unique." (Hawkins, 1978, p.110) In order to model the visible situation use, we need to represent the fact that the speaker's attention is focused at certain times on some objects, and that this focus of attention changes during a conversation [Grosz, 1977; Linde, 1979]. The plan discussed in (1.61) involves two boxcars, one in Elmira and one in Danville. In 29.5 the focus of attention is apparently Danville, since the reference to "the boxcar" is unambiguous even though three other boxcars are shown in the map. Yet, Danville clearly isn't the focus of attention during the whole dialogue, since another boxcar is discussed in 13.3-16.1, and at an moment in the discussion (do the manager and the system seem to perceive an ambiguity, not even when "that boxcar" is used in 13.5).

The definite descriptions "the boxcar" in 14.2 and "the boxcar" in 31.2 are cases of *anaphoric* use of definite NPs. According to Hawkins, we have an anaphoric use when the definite article is used to refer to an object explicitly "entered into the memory store" by a previous NP (Hawkins, 1978, p.106). The fragment in (1.61) illustrates another well-known fact about NP descriptions, namely, that when a definite description is used anaphorically, the only entities considered are those in the same discourse segment [Reichman, 1985; Grosz and Sidner, 1986; Fox, 1987].

Each felicitous use of a definite NP in a conversation that constitutes an apparent violation of the uniqueness (or identifiability) requirement on definite descriptions provides evidence that the participants to conversations are aware that information of different sorts is not mixed together in the common ground, but that different discourse topics (and subtopics) may be identified. In the TRANS conversations, for example, there are both statements about the world (as represented by the map) and statements about the plan (3.29) can be used both in response to (3.30a), in which case it is interpreted as an assertion about the current status of the world, and in response to (3.30b), in which case it would be interpreted as a statement about the state of affairs at a certain point during the execution of the plan.

- (3.29) B. Engine E1 is at Corning.
- (3.30) a. A. Where is engine E1?
- b. A. Where did we send engine E1?

In addition, different parts of the plan may be discussed at different times. For example, "the boxcar" in 31.2 is unambiguous, even though more than one boxcar has been mentioned as being part of the plan.

The topic of a discourse need not be an actual situation. For example, the participants in our conversations refer to objects and events which are part of the plan as if they were actual objects and events which actually occurred. In the following fragment, "a buscar" is introduced into the plan by the user in sentence 3.1, and then referred to in sentence 5.1, without the user specifying which buscar in the map he has in mind, if any.

- (3.31) 1 U okay, the problem is we better ship a buscar of oranges to  
Bath by 8 AM  
21 S okay  
31 U now, um, so we need to get a buscar to Comup  
where there are oranges  
32 there are oranges at Comup  
33 right?  
41 S right  
51 U so we need an engine to tow the buscar  
52 right?  
61 S right

Finally, (1.61) illustrates the need for interaction between the processes tracking attentional state and those performing intention recognition, recognized early on by Hobbs [1979]. Consider 3.1.2. For example, if the interpretation of the anaphoric definite "the buscar" were to take place before intention recognition has been performed, the discourse segment which includes 3.1.2 would not have been determined yet, hence all potential referents ought to be considered. Conversely, if intention recognition were to take place before the referent for "the buscar" had been identified, the plan reviewer ought to verify which action among all the actions involving buscars in the plan is being discussed. The most crucial contribution of Grosz and Sidner was to provide an hypothesis about how discourse segmentation and intensional structure might be related.<sup>32</sup>

### 3.7.2 The Common Ground as the Representation of a Discourse Situation

In [Poesio, 1993], I argue that there is a natural way to obtain a model of discourse that reflects the existence of multiple discourse topics and the effect of the focus of attention. The approach promised there, and incorporated with some revisions in CRT, is based on the perspective on context developed by Barwise and Perry [1983], according to whom the participants to a conversation make use of the knowledge that they are themselves part of a discourse situation in which they perform conversational actions. The common ground is interpreted in Conversation Representation Theory in a way that it is different from the way it is interpreted in DRT. The root DRS is not interpreted in CRT as a characterization of a situation being described by the discourse, but as a characterization of the discourse situation in other words, the conditions included in the root DRS describe facts that hold of the discourse situation, as opposed to facts that hold of the situation described by an utterance. In addition, the root DRS also contains information about the existence of a set of discourse topics, that are referred to by the utterances that are included in the discourse situation.

<sup>32</sup> An interesting phenomenon that I won't be able to discuss here is the fact that the participants in our conversations may refer to certain objects by means of descriptions which can only be interpreted if the hearer has "kept track" of the plan. For example, for definite the buscar that has oranges, in (1.3.3) or (1.61).

I show in [Poesio, 1993] that once definite description interpretation is formulated in situation-theoretic terms, the hypothesis that the common ground consists of information about the discourse situation gives us an explanation for a number of the observations about the pragmatics of definite descriptions reported in the previous section that does not require formal tools other than the notions of situation and of truth at a situation. And because of the way the semantics of the constructs of DRT is reinterpreted in CRT, this change in perspective does not lead to the introduction of hosts of new formal constructs, this will become evident in the next two sections.

For the purposes of this dissertation, it will be sufficient to consider a subset of the information contained in the discourse situation. In addition to information about discourse topics, we will need to assume that the common ground contains information about which utterances refer to which discourse topics. This information can be represented by assuming that the root DRS contains conditions that describe the utterance situations—or, as I will call them here, conversational events—that comprise the discourse situation. These conditions are important because I assume here the hypothesis discussed in [Poesio, 1993], according to which discourse segmentation is nothing else but the organization of conversational events in "threads", and therefore information about discourse segmentation can be reduced to information about particular situations that represent these threads. I will only use a very simple classification of conversational events: the conditions I use classify utterances as statements, questions, and imperatives. Conversational events and the representation of their impact on the common ground in CRT are discussed in §3.8.

### 3.7.3 Focus of Attention and Visual Attention

The model of the discourse situation must also include information about the current focus of attention. I reviewed in §3.7.1 Grosz's theory about the relation between focus of attention and the "visible situation" use of definite descriptions. According to Grosz, when an object is in the current *mutual* focus of attention, it can be felicitously referred to by means of a definite description even though other objects of the same type have been introduced in the discourse or are part of the visible situation.<sup>33</sup>

It is assumed within the Situation Theory literature that the object of visual attention is a situation; I make the same assumption, and call this situation *situation of attention*.<sup>34</sup> It is also assumed in Situation Theory that of all the objects in the visual field, only those within the current range of visual attention are actually "seen."<sup>35</sup> I therefore use a two-place SFE relation between an agent and a situation the agent is actively "looking at" analogous to the one used by Devlin [1991], ch 7). Finally, I use a relation *MOSE* between pairs of agents and situations to model the notion of current *mutual* situation of attention. Two agents *a* and *b* mutually see

<sup>33</sup> The "focus of attention" studied in the literature on definite description interpretation is called *object of visual attention* in some psychological work on visual attention [Alpert, 1987].

<sup>34</sup> The notion of *situation of attention* in Conversation Representation Theory is a role similar to Barwise and Perry's "object" we are attending to [1983, page 47]. Assuming that the focus of attention is a situation leads to simpler accounts relating the current focus of attention to the discourse situation and allows for more than one object to be in the current focus of attention.

<sup>35</sup> The distinction between "seeing" and "looking at" has been repeatedly discussed in the Situation Semantics literature [Barwise and Perry, 1983; Devlin 1991].

a situation  $s$ , written  $WSP(a,b,s)$ , if they mutually know that both of them see the situation. I assume that each situation of attention is shared.<sup>26</sup>

In the TRAINS conversations, the situation of attention is always a sub-situation of the 'map world' situation. The participants in our conversations do not, however, group the information from the map on the basis of some random order of selection; the sub-situations used as situations of attention always consist of the information about a town in the map at a certain point in time. Presumably, this is because the conversational participants only refer to situations they may expect the other participant to be able to 'build' as well.

I use the function  $PLACE(p,s,t)$  to denote the situation characterized by the facts which are true at location  $p$  at time  $t$  in situation  $s$ .

### 3.7.4 Common Ground and Mental States

What the model of discourse proposed in CRT does not include is a representation of the mental state of the participants to the conversation. Developing such an account would require solving an exceedingly large number of open problems, not the least of which is choosing a 'mental perspective' to model. The literature on discourse interpretation—Grosz and Sidner, Cohen, Levesque and Perrault, for example [Grosz and Sidner, 1986; Cohen and Levesque, 1990; Perrault, 1990]—aims at developing a *theory*'s logic, that is, a characterization of the agent's mental states and how utterances affect them 'from the outside': the motivation is that it is all we can get at, since all we have are transcripts of conversations. Ultimately, however, an agent's logic—that is, an account of what goes on in an agent's mind when bearing an utterance—should be most desirable for trying to understand the actual reasoning processes of people.<sup>27</sup> By basing my account on DRT, I have implicitly adopted an agent's perspective: the root DRS may be thought of as a characterization of one agent's beliefs about what's mutually believed. I have, however, simplified the task in a number of ways, not the last of which is the assumption that all information in the common ground is 'of the same type', namely, information about beliefs, as opposed to information about desires, goals, etc.<sup>28</sup> I have been able to make this simplifying assumption because I do not discuss speech act recognition, but it's clear that a more complete model would need to deal with this important aspect of discourse modeling as well.

## 3.8 DISCOURSE SEGMENTATION AND CONVERSATIONAL EVENTS

In [Pezzo, 1993], the connection between utterances and their 'discourse topics' is formulated in terms of a simple theory about the kind of information about speech acts that is part of the common ground. That proposal has been incorporated in CRT; I review it in this section.

<sup>26</sup> Note that in a logic of knowledge like SS this would follow from the definition of 'mutual seeing': the fact that each mutually seen situation is actual, and the veridicality axiom of Situation Theory.

<sup>27</sup> This aspect of discourse has been proposed by Devlin.

<sup>28</sup> A version of DRT that does not make this simplification, and in which every item in the common ground is marked with a color, belief, desire, want, etc. is sketched by Kamp in [Kamp, 1991].

### 3.8.1 Discourse Segmentation and Speech Acts

Grosz and Sidner [1986] account for the interaction between discourse segmentation and anaphoric uses of definite descriptions discussed in §3.7.1 by assuming that discourse segmentation is 'parasitic upon the intensional structure' ([1986], p. 180). According to G&S, whether a hearer interprets an utterance as being part of a discourse segment depends on whether the intensional(s) expressed by it (its *discourse segment purpose*) are related to the intensions expressed by the discourse segment.<sup>29</sup> Grosz and Sidner propose that intensions may be related in two different ways: when a discourse segment purpose is part of the satisfaction of another discourse segment purpose, the second purpose is said to *dominate* the first; if, instead, satisfying one intension is a prerequisite for satisfying a second one, the first intension is said to *satisfy* the second; and the second intension.

Grosz and Sidner's proposal effectively establishes a connection between discourse segmentation and speech act theory. In particular, identifying the segment in which an utterance is included becomes, in Grosz and Sidner's theory, a matter of inferring the connection between the intension(s) expressed by that utterance (essentially, the illocutionary act performed by the speaker of that utterance) and the current discourse intensions. Although many details still have to be worked out (in particular, no full account of illocutionary act inference exists), there is a widespread feeling that the basic ingredients of the account are correct.

My own proposals about discourse segmentation are an elaboration of Grosz and Sidner's. What I am concerned about is to show how Grosz and Sidner's proposals fit together with the idea that the common ground consists of information about the discourse situation in which the participants to a conversation find themselves. The connection is as follows: the participants to a conversation share not only knowledge that certain utterance situations have occurred, but also that (these utterance situations 'fit together' in courses of action). The organization of utterance situations in courses of action, and the dominance relations between courses of action, replace Grosz and Sidner's intensional structure. I discuss below how I propose to reconstruct their notion of focus space stack.

### 3.8.2 Speech Acts in Natural Language Processing

The aim of the work on intention recognition of Allen, Cohen, Levesque, Perrault and others [Cohen and Perrault, 1979; Allen and Perrault, 1980; Grosz and Sidner, 1986; Cohen and Levesque, 1990; Perrault, 1990] is to model the process by which the addressee comes to recognize the speaker's intensions in uttering a sentence. As the literal intension of an utterance can be rather different from the actual intension, recognizing these intensions may require complex reasoning.

The unifying characteristic of these models is the assumption that by uttering a sentence, a speaker is performing a speech act [Austin, 1962; Searle, 1969]. A declared goal of some recent research, most conspicuously the work of Cohen and Levesque [1990], is to derive the properties of speech acts observed in the earlier literature from general properties of actions, instead of stipulating the properties of illocutionary acts. Hence, part of the task of those engaged

<sup>29</sup> Actually Grosz and Sidner make a distinction between an utterance's intension(s) and the discourse segment purpose it carries. In simple cases like those discussed in this paper this distinction can be ignored.

in this line of research is to develop a model of actions from which the general properties of conversations may be derived.

The theories of action developed in this literature are, for the most part, of no concern to us, all we really need, for the purposes of this dissertation, is the assumption that the discourse situation includes actions performed by the participants to the conversation that are not in principle different from any other kind of action, together with an axiom in the effect that the occurrence of a conversational event causes a state of the hearer believing that a certain event occurred. An axiom of this kind is included in all theories of speech acts. Perrault [Perrault, 1990] calls it the *Observability Axiom* ([Perrault, 1990], p. 172)

*Observability* :  $\vdash DO_i \alpha \& DO_j \text{Obs}(s) \supset B_{j,t+1} DO_i \alpha$


$DO_i \alpha$  reads "x did  $\alpha$  at time  $t$ ," while  $B_{j,t+1}$  reads "y believes at time  $t$  that p." The axiom can be paraphrased as follows: if x performs action  $\alpha$  at time  $t$ , and y is observing x at time  $t$ , then y will believe at time  $t+1$  that x did  $\alpha$  at time  $t$ . As we will see below, the first step of interpretation in CRT is essentially a reformulation of the *Observability Axiom* in a DRT-like framework.

### 3.8.3 Speech Acts in Situation Theory and Conversation Representations Theory

As discussed in §2.2, in Situation Semantics, as well, it is assumed that discourse situation contains sub-situations, called 'utterance situations,' that correspond to speech acts. In other words, it is assumed in Situation Semantics that a discourse situation contains such sub-situations as the following utterance situation  $w$ , corresponding to the locutionary speech act performed by agent  $\alpha$  in telling to  $b$  that  $\phi$ :

$$|w| \models \text{TELL}(a,b,\phi)$$

Situation Theory's 'utterance situations' are called *conversational events* in CRT. The information that becomes part of the common ground as the result of a telling  $b$  that  $\phi$  includes information about the occurrence of a locutionary act, this effect can be represented by adding to the next DRS representing the common ground a new discourse marker  $ce$  and an event description characterizing  $ce$  as an instance of the following improved DRS as the result of an utterance



The event description reads: the (conversational) event  $ce$  can be completely characterized as an instance of the type of event in which the predicate  $\text{TELL}(a,b,s,K)$  holds, where  $s$  is a situation (description) stating that a certain situation  $s$  is of the type specified by the content of the locutionary act, characterized by the DRS  $[\phi]$  containing  $\phi$  as side condition.

The situation  $s$  is called in Situation Theory the *described situation*. As it's not the case that all utterances describe situations, and assertions are not always about actual situations, I use instead the term *discourse topic* to refer to the situation or situation kind that an utterance is 'about'.

A set of *conversational event generation rules* (CEGR) specify what gets added to the root DRS as the result of various types of locutionary acts. The conversational event generation rule for declarative sentences can be rephrased as stating that the occurrence of a (conversational) event  $ce$  of *speaking* hearing that  $\phi$  results in augmenting the common ground (i.e., the root DRS) with information to the effect that  $\phi$  is true. The connection between this rule and Perrault's observability axiom should be obvious: rules in more detail below, after having discussed discourse segmentation.

As I intend to ignore in this dissertation the problem of inferring illocutionary acts,<sup>40</sup> I made the following simplifying assumptions. I assume, first of all, that the locutionary act associated with an utterance can be recognized on the basis of syntactic and prosodic information only, and that the conversational event generation rules originate locutionary acts, as in Hwang and Schubert's proposal [Hwang and Schubert, 1993]. I distinguish between three classes of locutionary acts: *TELL*, *ASK*, and *INSTRUCT*.<sup>41</sup> Finally, and most importantly, I assume that, although Grosz and Sidner's intentional structure depends on the relation between the *illocutionary* acts associated with utterances, one can derive relations between the locutionary acts from the relations between illocutionary acts. In other words, I assume that, if the locutionary act  $ce$  of saying that  $p$  generates the illocutionary act  $i$ , and if  $i$  is inferred to be subordinate to the discourse segment purpose associated with the segment  $ds$ ,  $ce$  will also be subordinate to  $ds$ . I leave for another occasion the discussion of the inferences that result in associating the illocutionary act  $i$  with the locutionary act  $ce$ , and in establishing that  $i$  is subordinate to  $ds$ .

### 3.8.4 Conversational Threads, Discourse Topics, and Discourse Segmentation

The effect of discourse segmentation on anaphoric accessibility is formalized by Grosz and Sidner by means of an abstract data structure called the *focus space stack*. As long as an utterance is part of the current discourse segment, the discourse referents evoked by that utterance are added to the 'focus space' on top of the stack, and the discourse referents in that focus space are accessible for anaphoric reference. When an utterance introduces a discourse segment subordinate to the current one, a new focus space is pushed onto the stack. When an utterance completes the current discourse segment, the current focus space is popped from the stack.

I believe that Grosz and Sidner's proposal can be recast in a straightforward way using situations and their relations instead of focus spaces. Moreover, I believe that by reformulating their account in this way, a better understanding of what the theory predicts is gained.

I propose, first of all, that just like all other events, conversational events are parts of courses of actions. This notion was introduced by Barwise and Perry in [1983] without much discussion.


<sup>40</sup>The approach to speech act interpretation taken in the TRAINS system is briefly discussed in [Traum et al., 1994].

<sup>41</sup>These locutionary acts roughly correspond to the 'surface speech acts' used in the Artificial Intelligence literature on speech act interpretation, e.g. [Allen and Perrault, 1990].



what I mean by a course of action here is a collection of events that can however be referred to as a unique object, analogously to a group in Link's (1983) sense. A total order relation is defined on courses of action, so that we can talk about the last event in a course of action, the previous event, and so forth. (See Chapter 4.) I use the term *conversational threads* for courses of actions involving conversational events. Conversational threads are the CRT equivalent of discourse segments.

Secondly, I propose that the information added to the common ground as the result of an utterance includes not only the information that a conversational event occurred, but also that that conversational event is part of a conversational thread yet to be determined. This is represented by the following CRT expression



Note that two parameters are used in this situation description. One parameter,  $\gamma_1$ , represents the conversational thread. One of the tasks of discourse interpretation is to anchor this parameter (i.e. to identify the conversational thread of which a conversational event is a constituent. We won't be concerned with this aspect of discourse interpretation, in general, whether a conversational event is perceived as part of a conversational thread depends on how the intentions expressed by that event are related to the intentions expressed by  $\gamma_1$  at thread

The second parameter in (3.32),  $\alpha$ , represents the discourse topic that the utterance is about. It may be identified as well. What makes a set of conversational events a thread is the fact that they are all about the same discourse topic. The specific events described by the conversational events in a thread are sub-situations of the discourse topic of that thread. The discourse topic of a conversational thread  $c$  that dominates the conversational thread  $c'$  is a sub-situation of the discourse topic of  $c'$ .

If we treat discourse topics as situations, we get a natural reformulation of the focus space (pushing/popping) model. The properties that Grosz and Sidner attribute to the focus space stack are also the properties of situations. Each operation on the focus space stack can be reformulated as an operation on situations. 'adding to a focus space' corresponds to 'adding new constituents', 'pushing' corresponds to 'create a new situation which informationally includes the previous situation', and 'popping' corresponds to 'selecting a situation informationally included in the previous one'. For each NP of the current utterance which introduces an anaphoric antecedent (and hence a new constituent in the addition of an object to the focus space), a new constituent is added to the current situation. An utterance that opens a sub-segment results in a new conversational thread being open, whose discourse topic informationally subsumes the discourse topic of the previous conversational thread. An utterance that closes a discourse segment and pops to a previously opened discourse segment results in the corresponding discourse topic becoming current.

Replacing the focus space stack with a structure of situations has several advantages. First of all, it is much simpler to understand the connection between this proposal and Grosz's earlier work on implicit focus (Grosz, 1977). The objects in 'implicit focus' are simply the constituents of a discourse topic not explicitly mentioned in the conversation. (For example, the objects

needed to ensure that the plan can get carried on: thus, after referring to "an engine", we can refer to "the driver", and so forth.)

Secondly, we get a principled way to choose the discourse segment that the referent of a NP is part of in case nested segments are present, choosing the 'top of the stack' - often results in undesirable pops. For example, imagine that we plan to move a boxcar of oranges from Bath to Cornwall, and we start talking about a sub-plan: say, moving a boxcar from Avon to Bath. Say that a subject that gets used throughout the plan (e.g., the engine moving the boxcar) is first mentioned during this phase of the conversation. According to the focus space approach, we ought to make the engine part of the focus space associated with the sub-plan. But this would predict that when we have completed the sub-plan, the engine should no longer be accessible, which is false. We can clearly refer to that very same engine when proceeding to a subsequent part of the plan (for example, sending the boxcar full of oranges to Cornwall). If we assume that we are talking about *situations*, and we know something about the kind of situation we are trying to achieve (say, we know that an engine may be involved in more than one part of the plan), we may hypothesize when the engine is first introduced whether this is going to be a 'local' constituent, or a 'global' constituent of the plan.

A final comment concerns the separation made here between discourse segmentation and the focus of attention, that in Grosz and Sidner's theory are merged.<sup>43</sup> GAS imply that their focus space stack is a formalization of all the relevant aspects of the attentional state. They lack background knowledge, however, the need for additional mechanisms such as centering for the purpose of modeling pronoun resolution (p.191) and do not make specific suggestions concerning the segmentation of focus of attention and discourse segmentation. I believe that maintaining a distinction between the structures used to interpret anaphoric and visible situation uses of definite descriptions leads to a clearer theory.<sup>44</sup>

### 3.8.5 Conversational Event Generation Rules

The first step of the version of the DRT construction algorithm discussed by K&R, applied to sentence 3, consists of adding the s-structure of the sentence to the root DRS. As discussed above, the first step of the KRT construction algorithm is to add to the root DRS (that we now use as providing a characterization of the discourse situation) an underspecified representation characterizing the conversational event that just occurred. Conversational event generation rules specify what gets added to the common ground, according to the kind of conversational event

In this section I discuss the three CEGRs used for the TRAINS conversations. For simplicity, these rules are presented using the format adopted by Kamp and Reyle to present their operations on condition sets; the conversational event generation rules are not, however, rules like the other two. The first two CEGRs are the root DRSs of the TRAINS conversations. That is, they are not rules that are applied to a logical form once this logical form gets added to the root DRS. The CEGRs presented below are the best seen as summarizing the contribution to the common ground of an utterance, between the its syntactic structure and prosodic features.

<sup>27</sup> Discourse topics provide the segmentation mechanism necessary for dealing with anaphoric uses of definite descriptions: my model of the focus of attention is discussed in §6.1.


<sup>43</sup> Some recent work on attention also points out the need for a separation between the attentional components involved with different scenes (Allport 1987).

I assume, following Hwang and Schubert [Hwang, 1992, Hwang and Schubert, 1993], that the logical form includes a **surface speech act operator** DECL for declarative utterances, QUES for interrogatives, and IMPER for imperative utterances. I only assume the existence of such operators, without providing a truth-conditional characterization of the distinction between declarative sentences, interrogatives, and imperatives.<sup>44</sup>

I also assume that these operators are the translation of syntactic markers that occupy at structure the position of [Spec, CP], the 'highest' maximal projection contained in a clause. According to the analysis of wh questions adopted in Government and Binding theory as presented by Haegeman [1991], the wh-phrases occupy at structure the position of specifier of CP. The sentence 'Which engine shall we send to China?' for example, is analyzed as in (3.33): the auxiliary 'shall' moves to the position of head of CP leaving behind a trace, while the wh phrase moves to the position of specifier of CP.



The presence of a wh phrase in [Spec CP] indicates that the sentence is a wh question. We can generalize this proposal and assume that the surface speech act expressed by a sentence is always specified by an operator in [Spec CP]. In the case of wh questions, the operator is the translation of the wh phrase. In the case of other kinds of sentences, I assume that the position of [Spec CP] is occupied by one of the three surface speech act operators listed above. For example, the declarative sentence 'John loves Mary' results in a logical form that, ignoring the analysis of the clause 'John loves Mary', is as follows:




I emphasize that nothing in the theory depends on the surface speech act indicators assuming this particular syntactic position.

The first conversational event generation rule, CEGR DECL, specifies what gets added to the common ground in CRT as the result of a declarative utterance. Declarative utterances are utterances whose logical form includes a DECL operator in the specifier position of CP. As discussed above, uttering a declarative sentence results in adding to the common ground a new conversational event of type TELL. *ce* I also mentioned that TELL is a relation between two


<sup>44</sup>Classical references on interrogatives are [Beck, 1982, Berman, 1991, Engdahl, 1986, Ginzburg, 1992, Groenendijk and Stokhof, 1990, Karttunen, 1977]. The literature on imperatives is much more limited: two prominent examples are [Hilborn, 1982, Mew, 1991]. A computational treatment of a class of imperatives is discussed in [Hilborn, 1993].

agents and a situation description. According to CEGR DECL, the contribution to the common ground of a declarative utterance is an instance of a TELL event whose content is a parametric situation description that includes as a condition the subree of the logical form that appears in [Comp, CP].<sup>45</sup> The new conversational event is in turn part of a conversational thread (discourse segment) to be identified




The second conversational event generation rule, CEGR QUES, describes the effect on the common ground of a yes-no question, defined as a sentence whose logical form includes a QUES mood operator in [Spec, CP]. Such an utterance results in the addition to the common ground of a surface conversational event of type ASK.

<sup>45</sup>In general, every CP in a sentence results in a new situation description.



Finally, the rule **CEGR IMPER** specifies the effect on the common ground of uttering an imperative sentence: an event description characterizing a conversational event of type **INSTRUCT** is added to the root DRS. Note that this type of conversational event differs from the previous two in that its third argument is a situation *kind*, as opposed to a situation. This captures the intuition that imperative sentences are tenseless, as well as the fact that the complement of "instruct" is an infinitival clause.<sup>46</sup>

<sup>46</sup> The treatment of infinitives in CRT is derived from Huang and Schieffelin, where infinitivals denote situation kinds. Their treatment in turn is a modification of Chierchia's [1994], where infinitivals are taken to denote properties. I discuss the semantics of infinitivals in §5.8.



### 3.8.6 Micro Conversational Events

The main reason I mentioned above for assuming that the common ground includes information about the occurrence of conversational events (instead of only information about the propositional content of these events) was the need to explain how a conversation may have more than one discourse topic. But when one starts looking at actual conversations, one soon realizes that there are additional reasons for taking such a step. One of these reasons is the fact that single lexical constituents of a sentence, and perhaps even the phonemes that form these lexical items, appear to have an impact on discourse interpretation possibly separated from the impact of the sentence as a whole, thus their occurrence must be recorded in the common ground. In this section I briefly discuss this point, and I sketch a proposal about discourse interpretation that accounts for these data, a more detailed discussion is in preparation.

Two interpretation processes that appear to depend on the occurrence of such utterance fragments in the common ground are grounding and repair. Grounding [Clark and Brennan, 1990, Clark and Schaefer, 1989] is the process by which speaker and hearer achieve 'mutual understanding'. Typically, a speaker does not go on with an utterance until she's finished, but waits for the hearer to acknowledge parts of what she just said before continuing. Consider the following example, a slightly edited fragment of one of our transcripts:

- (3.34) 11 U We're gonna take engine E2  
 21 S okay  
 31 U hook it to the boxcar at Elmira.  
 41 S okay  
 51 U and take it back to Elmira

If we assume that each contribution to the conversation is recorded in the common ground, we must assume that the fragments in 1, 3, 1 and 5.1 are added to the common ground, and then somehow "merged" into a complete utterance. If we assume that the common ground is a picture of the discourse situation, and that the effect of 1.1 and 3.1 is to add *micro conversational* events to the common ground, the fragment in (3.34) does not present particular problems. Also, the reasoning involved in assembling 1, 3, 1 and 5.1 into a single conversational event is in many ways similar to the reasoning involved in assembling several conversational events into a single conversational thread.

As neither 1.1 nor 3.1 denote sentential objects (thus, these objects cannot be simply added to the current DRS), the only alternatives to the "micro conversational events" hypothesis are to assume either that grounding takes place prior to semantic interpretation, or that each utterance fragment is "completed" obtaining an object of type *t* before adding it to the common ground. Because an utterance fragment can be completed in many different ways, the second proposal can only be made to work if the type *t* expression added to the common ground is somehow underspecified, so this proposal can be seen as a simplified version of one variant of the micro conversational events theory. As far as the first alternative is concerned, the phenomenon of repair shows that it won't work. For example, the transcript in Fig. 12 includes the following fragment:

- (3.35) 9.1 U we should arrive the engine at Avon  
 9.2 engine E1 in (inc)  
 10.1 S engine E1  
 11.1 U E1  
 12.1 S okay  
 13.1 U engine E1, so back to (inc)  
 13.2 or we could actually move it to Danville, and pick up  
 the boxcar there  
 14.1 S okay

That the "system" performs a repair in 10.1 indicates that the referent of "the engine at Avon" in 9.1 has been interpreted, and found to be different from the referent of "engine E". This indicates that 9.1 and 9.2 have already been interpreted at the point in time in which the system utters 10.1, which is prior to the user's completing its sentence.

There is an additional reason to assume that the common ground contains utterance fragments— in fact, that utterance fragments may be taken to be part of the common ground no matter what kind of discourse is processed, including written text. This is the fact that listeners may have different preferred readings for sentences whose structure is different, but that arguably are semantically ambiguous in the same way (that is, they denote the same "set of propositions"). One example is the contrast between active and passive sentences discussed by Kurtzman and MacDonald

- (3.36) a. Every kid climbed a tree.  
 b. A tree was climbed by every kid

Assuming a theory of the passive such as Parsons' [1990], both the logical form that translates (3.36a) and the one that translates (3.36b) denote the same set of interpretations if we use the storage technique to compute the derivation, yet we have seen how people have very different preferences when interpreting the two sentences. The same is true for the examples discussed by Reinhart, in which, again, "two sentences that truth-conditionally appear to be equivalent, result in very different scopings" (cited).

- (3.37) a. Some reporter interviewed Kissinger in every town  
 b. In every town, some reporter interviewed Kissinger

We could give two explanations for this: either it is the case that discourse interpretation works off some sort of representation, so that it is sensitive to differences in representation even when they do not correspond to differences in informational content (one often finds this idea, albeit implicitly, in the DRT literature), or else it must be the case that the common ground resulting from an utterance of (3.37a) is different from the common ground resulting from an utterance of (3.37b), but this difference has to do with the discourse situation, rather than with the semantics of the sentence. I propose that the latter is the case, and that the additional information consists of information about the occurrence of utterance fragments and how these fragments compose into larger fragments. For example, (3.37a) differs from (3.37b) because the latter corresponds to an utterance fragment in which the temporally ordered fragments "in every town" and "some reporter interviewed Kissinger" are composed.<sup>47</sup>

More precisely, I claim that after hearing the PP "in every town", the common ground is augmented with a new object, that I call utterance fragment. I assume a function *U* that maps logical form arguments into properties that are true of utterance fragments whose semantic translation is an instance of that logical form fragment. For example, hearing the PP "in every town" results in adding to the common ground a new discourse marker *f<sub>1</sub>*, and the following condition

$$\begin{array}{c} \text{PP} \\ \text{U} \lambda P \lambda y \forall N x [x \models \text{TOWN}(x)] \\ \text{([N](P)(y))} \end{array} \quad M(f_1)$$

An utterance fragment gets also added to the common ground upon hearing the TP "some reporter interviewed Kissinger", let's call this *f<sub>2</sub>*. In fact, I propose that the common ground contains one such utterance fragment for each lexical item in the utterance, together with utterance fragments that correspond to the sentence constituents that are built out of more basic constituents, whenever the listener believes that the speaker intended the utterance fragments to be "put together" in that fashion. I propose to represent this by allowing logical forms to have

<sup>47</sup>The repeated occurrence of the definite "the latter" in the previous paragraph is not entirely fortuitous. I believe that phrases such as "the former" and "the latter" provide additional evidence that the common ground of a conversation (or of a text) includes information about which noun phrases were used in which temporal order, rather than simply about the derivation of sentences.

utterance fragments as nodes, thus, the utterance fragment obtained by combining the utterance fragment for the PP "In every town" with the TP "some reporter interviewed Kissinger" is as follows:



A logical form one of whose nodes is an utterance fragment  $f$  has the same denotation as the logical form in which  $f$  has been replaced by the logical form fragment  $f$  such that  $\|f\|$  is the case.

I furthermore, I propose that the common ground contains, for each utterance fragment *f*, a micro conversational event of uttering *f*. Thus when hearing 'in every town', the following is also added to the common ground:


UTTER(*spr*, *s*<sub>1</sub>)

(One such micro conversational event is also added to the common ground for each utterance fragment that gets assembled

In general, I propose that utterances are processed as follows. A listener begins to process a contribution to a conversation by adding to the common ground a new course of action. For example, when beginning to process the sentence 'A boxcar left', i.e., when starting to process the first lexical item, 'the', a conversational participant adds to the common ground a situation description of the firm.

$$\begin{array}{c}
 \begin{array}{c}
 \text{Det} \\
 \text{u} \\
 \lambda P \lambda Q [Q(d) \wedge P(d)]
 \end{array}
 \quad
 \begin{array}{c}
 K(x) \\
 \begin{array}{c}
 \text{[INTERPRET } f, i] \\
 u_1 \quad @ \quad f_1 \\
 u_1 \quad \sqsubseteq \quad i_1
 \end{array}
 \end{array}
 \end{array}$$

Upon hearing the second lexical item, the noun "boxcar", a new utterance fragment and a new micro conversational event are added to the description of the situation. Once the listener recognizes that the determiner and the noun are meant to constitute an NP, a new utterance fragment and conversational event are similarly added. At any moment during this processing a repair may be needed, and in fact, it's by no means certain that the contribution will be concluded and result in a conversational event, if that happens, however, the conversational event is also guided to the common ground, together with the information that it was generated by the micro conversational events. We obtain, in other words, something like the following:



This solution has several advantages. First of all, the two 'Kisinger' sentences discussed by Reinhart, although truth-conditionally equivalent, result in different additions to the common ground (i.e., different micro-conversational events and utterance fragments are added in the two cases). Secondly, it provides a level of interpretation at which grounding and repair phenomena can be dealt with without assuming that they take place prior to interpretation (3.35), for example, can be analyzed as a case in which the interpretation of the definite description 'the engine at Avon' is 'triggered' by an utterance fragment, thus takes place before the rest of the sentence is interpreted.

In the rest of the dissertation, I assume that this information is added to the common ground, without however presenting it. I also usually write the rules that depend on certain syntactic configurations without specifying that they actually depend on certain utterance events having occurred, because that's always the case.

### 3.9 WEAK AND STRONG DEFAULTS IN DISCOURSE INTERPRETATION

The construction rules introduced in DRT can be seen as a generalization of the notion of inference rule used in traditional logic. The discourse interpretation rules discussed in Chapter 6 encode both commonsense reasoning and the 'interpretive' reasoning involved in assigning a scope to certain operators or in identifying the interpretation of certain contextually-dependent expressions. Some discourse interpretation rules resemble the rules of the DRS construction algorithm in that they are 'triggered' by particular syntactic configurations, in combination with facts about the discourse situation (I will reserve the term *model construction rules* to indicate the rules that more closely correspond to the rules of the DRS construction algorithm, modulo their dependency on discourse interpretation formulated as the Anti-Random condition), whereas other rules are more like the axioms normally used to formalize commonsense reasoning.

As discussed in §1.5, I have also adopted the hypothesis that there are two basic categories of construction rules, *strong* interpretation rules (also called *signal forced*), and *weak* default rules (or *expectation-based*).<sup>40</sup> This hypothesis corresponds to the hypothesis that, when interpreting an utterance, people rely on expectations that however can be overridden by 'signals' such as lexical priming or anaphoric associations, independently arrived at in work on Conversational Analysis (Levinson, 1983) definite description interpretation (Haviland and Clark 1974, Crain and Steedman, 1985), pronoun interpretation, and tense interpretation (Kanevama *et al.*, 1993).<sup>41</sup>

I have proposed in §1.5 that the same strategy is at work in scope interpretation, and I have mentioned there that this interaction can explain the scope preferences observed in the literature and in our conversations. I proposed there that the scope preference for Subjects to be interpreted as part of what's given is a 'weak' interpretation rule that can be overridden by other discourse interpretation procedures—for example, the procedures that interpret definite descriptions or those that assign to preposed PPs the role of given—all of which are, I propose, strong interpretation procedures. This accounts, I propose, for the contrast between (1.40a) and (1.40b).

- (1.40) a Some reporters interview Kissinger in every town  
b In every town, some reporters interview Kissinger

Conflicts arise whenever two rules of the same 'strength' may apply: for example, in the case of passives, where two 'weak' rules generate a conflict.

- (1.41) A tree was climbed by every kid


In general, computing the extension of a theory with potentially conflicting rules is not an easy task. To keep matters simple, I have assumed a conflict resolution strategy based on several phases of rule application. During the first phase, only the strong rules may apply. This may

<sup>40</sup> This distinction does not apply to model construction rules, but to the construction rules that model discourse interpretation.

<sup>41</sup> Hest (Hest 1987) p. 190 reports work by Marden Wilson and Tyler that also suggests the priority of bottom-up processes (Wilson, Nelson and Tyler 1980; Tyler and Marden Wilson 1982).

result in a partial or perhaps even complete hypothesis about the utterance. When no strong rule may be applied anymore, a second phase ensues, during which only the weak rules may apply.

To illustrate, consider (3.40b). During the first phase, in which strong defaults are applied, the relevant set of towns is identified as the topic because of its occurrence in a quantified NP embedded in a preposed PP (see §6.2), and the model construction rule for universal quantifiers applies. This results in an hypothesis that can be schematically represented as follows:



After the closure of the hypothesis under the strong defaults has been computed, the second phase begins, and weak interpretation rules may apply. At this point, however, the sentence topic has already been determined, as well as the relative scope of the NPs 'every town' and 'some reporter'. In other words, the strong default has overridden the weak default.<sup>42</sup>

As said above, the result of the discourse interpretation procedure is, in general, a set of hypotheses, each representing a distinct, consistent and complete interpretation of the utterance. In the case of a scopally ambiguous sentence, for example, one hypothesis will be produced for each interpretation of the sentence that the current context 'suggests'.

### 3.10 SUMMARY: A REVISED MODEL OF DISCOURSE INTERPRETATION

Let me summarize what I have proposed in this chapter. I have proposed to replace the DRS construction algorithm with a discourse interpretation procedure that maintains many characteristics of the procedure developed by Kamp and colleagues, but is based on the ideas (introduced in Situation Semantics) that common ground is best seen as a representation of the discourse situation and as such it contains information that is pragmatic in nature, such as information about the occurrence of conversational events and their temporal order, or about the attentional aspects of the discourse situation. This position will hardly seem surprising to those familiar with the literature on discourse interpretation (Hobbs 1979, Grosz and S. Jeter, 1986, Clark, 1993), but is by no means commonly accepted in the formal semantics literature. I motivated this shift in perspective by looking in §1.7 at the kind of information that seems to be used to interpret definite descriptions.

The CRT construction algorithm is still based on the idea that the content of the common ground can be represented in terms of structures consisting of sets of markers and sets of

<sup>42</sup> An interesting problem for this proposal is represented by sentences such as

- (1.43) Every kid climbed the tree

The definite NP 'the tree' appears to take wide scope even when the sentence is processed without any previous context. One could take the route that when a definite is encountered in solution, a unique referent is simply made up. This however makes it difficult to explain why (1.44) is well-formed.

- (1.44) John dated the woman

conditions, and that the 'update potential' of an utterance is best formulated in terms of operations that replace partially underspecified interpretations with others which represent more specific hypotheses. I argued that both syntactic and lexical information need to be taken into account during discourse interpretation, thus the logical forms introduced in §1.5 are a better characterization of the input to discourse interpretation than the purely syntactic 'uninterpreted conditions' assumed in the version of the algorithm presented by Kamp and Reyle. An independent motivation for introducing these logical forms was mentioned in §1.5, where they were introduced to characterize the distinction between semantic ambiguity and perceived ambiguity.

The input to the CRT interpretation procedure is specified by **conversational event generation rules** that add to the common ground a set of **utterance events** and, possibly, a **surface conversational event** generated by the utterance events. The surface conversational event is an underspecified characterization of the current utterance. The result of adding the surface conversational event and the utterance events to the most DRS results in a condition set that I call **preliminary hypothesis**. This hypothesis is the starting point of the discourse interpretation procedure proper. Discourse interpretation is seen in CRT as the process by which parameters get anchored. As soon as a parameter gets anchored, the common ground may be updated, which results in a more specific form. During discourse interpretation, new hypotheses are obtained by applying **discourse interpretation rules** to existing hypotheses. Because mutually incompatible discourse interpretation rules may be 'triggered' by the same hypothesis, multiple hypotheses may in general be obtained, each hypothesis corresponds to one of the extensions of the inference system described in §4.4. The result of the discourse interpretation procedure is a set of alternative hypotheses about the kind of speech act performed with the utterance. It is then the job of discourse reasoning and of the planning components either to choose the most plausible hypothesis, or to decide that the ambiguity cannot be resolved and therefore apply whatever ambiguity resolution strategy is deemed appropriate.

The algorithm I adopt also differs from the DRS construction algorithm in that it takes into account the defeasible nature of most of the interpretive inferences, and therefore its result is a set of hypotheses about the utterance; a simple conflict resolution procedure is adopted that assumes the existence of an 'external' evaluation function (e.g., provided by a plan reasoner). I have sketchily described how an inference system may be defined that would allow us to treat the DRS construction rules as inference rules on par with the rest of the rules used for discourse interpretation, thus allowing for a smooth integration of various interpretation processes. The inference system is described in greater detail in §4.4.

## 4 Formal Aspects of CRT

In this chapter I discuss the formal aspects of Conversation Representation Theory in some detail. Although CRT cannot be called a true logic yet—for example, the notion of consequence is still in large part unexplained—I provide here a specification of the denotation of the expressions of CRT detailed enough to show how the formalism does the job. The chapter can be skipped by those who are not interested in the formal details.

### 4.1 TYPES AND THE MODEL

#### 4.1.1 Types in CRT

The semantics of CRT is based on a set  $T$  of semantic types, the smallest set such that

1.  $e$  is a type.
2.  $s$  is a type.
3.  $t$  is a type.
4. If  $\tau$  and  $\tau'$  are types,  $\langle \tau, \tau' \rangle$  is a type.

The set of basic types includes a type  $s$ —the type of situations—that can be freely used to obtain functional types, unlike type  $s$  in Intensional Logic. The set of meaningful expressions of type  $\alpha$  is indicated by  $ME_\alpha$ . The set of non-logical constant expressions of type  $\alpha$  is indicated as  $CE_\alpha \subseteq ME_\alpha$ .

In Intensional Logic, the **evaluation function**  $\| \cdot \|^{M,g,w,t}$  provides the denotation of expressions with respect to a model, an assignment, a world  $w$ , and a time  $t$ , the type assigned to (natural language) sentences in Montague Grammar is  $t$ . An equivalent way to assign a semantics to expressions is to 'eliminate' the world parameter from the denotation function, and take objects of type  $t$  to denote total functions from possible worlds to truth values.<sup>1</sup> 'Conservative' versions of Situation Theory such as Episodic Logic generalize this second approach to the definition of

<sup>1</sup> This alternative way of formulating a semantics for logical expressions is explicitly adopted, for example by Crestwell in [1971].

the evaluation function in two directions (the translations of) sentences denote functions from situations to truth values, and these functions are allowed to be *partial*. Thus, in Epistemic Logic, all expressions, including constants, denote partial functions from situations to their values in Intensional Logic.<sup>2</sup>

In CRT, a further generalization is introduced. Having hypothesized that ambiguous sentences denote sets of functions from situations to truth values, the simplest way to include in the language of CRT expressions used to translate ambiguous sentences, and to allow for these expressions to be conjoined with other sentence-level expressions (i.e., expressions of type  $t$ ) is to have all expressions of type  $t$  denote such sets. The same kind of reasoning applies for the denotation assigned to expressions used to translate arguments, since we have to allow for semantically ambiguous arguments, such as pronouns. Again, starting from the 'typing' used in Epistemic Logic, where arguments denote functions from situations to objects in the universe of discourse, one can generalize that type to sets of functions from situations to objects, so that the same type gets assigned both to pronouns (that are semantically ambiguous) and to proper names.

One way to do this would be to assign to natural language expressions types of the form  $(\alpha, \beta)$ , where  $\alpha$  is the type these expressions have in Epistemic Logic.<sup>3</sup> The problem is that with this type assignment, whether an expression is ambiguous or not depends entirely on the model, and the assignment of function, which is not what we want: what we want is a logic in which certain expressions indicate the existence of ambiguity, as it were. This problem is best illustrated by considering the case of variables vs. parameters. What we want is for the parameter used to interpret the pronoun 'he' in 'He climbed a tree' to indicate ambiguity, hence to denote a set of functions from situations to objects of cardinality greater than 1, whereas, say, the variables occurring in the interpretation of quantified statements like 'Every kid climbed a tree' that are not ambiguous, should denote singleton sets of such functions. But if we use  $(\alpha, \beta)$  as the type of both variables and parameters, there is no way to enforce this distinction other than adding a constraint to the effect that all assignments must assign singleton sets to variables of type  $(\alpha, \beta)$ .

I followed, therefore, a different route, namely, making all CRT expressions set-valued, and defining the evaluation function  $|||$  in such a way as to have it assign a set of values to each meaningful expression, a singleton set to unambiguous expressions, a set of cardinality greater than one to ambiguous ones. If we do things this way, and maintain the 'situated' approach to evaluation found in Epistemic Logic, we end up assigning to natural language expressions in Conversation Representation Theory the same types that these expressions have in [Dowty & Conversation Representation Theory] (as revised by Partee and Rooth [1983]), with the difference that, when I talk about 'meaningful expressions of type  $\tau$ ' below, therefore, I am really talking about expressions that denote sets of functions from  $S$  to elements of  $D_\tau$  (the domain of type  $\tau$ ). This is illustrated in the following list:

- The denotation of sentences,  $r$  in Dowty, Wall and Peters' system and  $r$  in CRT. However, the evaluation function of CRT assigns sets of functions from situations to truth values as the denotation of objects of type  $t$ .

<sup>2</sup>Let us ignore here the further complication that Epistemic Logic is not based on types.

<sup>3</sup>That  $\tau$  if the denotations used in Epistemic Logic were to be formulated in terms of types. Also, as just mentioned, an object of type  $\alpha$  would denote a function from situations to the type of objects that  $\alpha$  denotes in Intensional Logic.

- The denotation of proper names,  $e$  in Dowty, Wall and Peters' system (as revised by Partee and Rooth [Partee and Rooth, 1983]), and  $e$  in CRT as well. This means that proper names denote (sets of) functions from situations to elements of the domain in CRT.
- Properties:  $(\alpha, \beta)$  in Dowty, Wall and Peters, and  $(\alpha, \beta)$  in CRT as well, but in CRT, meaningful expressions of type  $(\alpha, \beta)$  denote sets of functions from  $S$  to  $D_\beta(\alpha, \beta)$ .
- Relations: same as properties  $(\alpha, \beta)$  in Dowty, Wall and Peters and  $(\alpha, \beta)$  in CRT but in CRT meaningful expressions of type  $(\alpha, \beta)$  denote sets of functions from  $S$  to  $D_\beta(\alpha, \beta)$ .

#### 4.1.2 Objects and Their Structure

Meaningful expressions are assigned a value with respect to a universe  $U$ . I use below the notation  $\alpha$  to indicate that  $\alpha$  'stands for' an object in  $U$ , i.e., it is part of the metalanguage, as opposed to being a meaningful expression of the object language.

Following Link [1981], I assume that the universe of discourse  $U$  is 'structured', in the sense that it contains, in addition to 'atomic objects' groups of these objects used to interpret plurals, and that groups and atomic objects are organized into a semi-lattice.<sup>4</sup> I assume in the following that  $U$  is such a semi-lattice.

I also borrowed some notational devices from Link. I use the  $P$  notation to indicate predicates whose denotation are non-atomic objects built up from atomic objects in the denotation of  $P$ . I also use Link's logical predicate  $ATOM(x)$ , that is true of a term  $x$  if  $x$ 's denotation is one of the 'basic' elements of the universe (those that have no sub-elements).

$U$  includes a set  $K$  of kinds as a subset; an isomorphism  $\lambda$  is defined from propositions to situation kinds.

#### 4.1.3 Situations

The models with respect to which a CRT expression is evaluated include a set  $S$  of situations. Two partial order relations are defined over the elements of this set. There is, first of all, a relation modeling 'real inclusion' between situations in a world: this reflects the fact that a domain situation—the situation representing a particular event of Kim buying popcorn, say—may be part of another domain situation, say, Kim going to the movies yesterday night. This notion of inclusion between situations is modeled by assuming that a semi-lattice is defined over the set of situations  $S$ , the meta-language relation  $\sqsubseteq$  (and its object-language correspondent  $\sqsubseteq_o$ ) are used to talk about 'this kind of situation inclusion'.<sup>5</sup>

The second notion of inclusion is defined over *information*, that is, between *abstract* situations (situation types). An abstract situation is included in another abstract situation if every

<sup>4</sup>A semi-lattice is a structure on which a partial order relation is defined and which includes a 'top' or maximum element, although it does not include a single bottom.

<sup>5</sup>More precisely,  $\sqsubseteq$  is a function from predicates to predicates, that maps a predicate  $P$  into the predicate  $P^*$  whose extension includes groups of objects in the extension of  $P$ .

<sup>6</sup>This relation is similar to the part-of relation  $\prec$ ,  $\subset$ ,  $\supset$  used by Barwise and Perry [1983].



(domain) situation in the set denoted by the fix is also in the set denoted by the second. This relation between abstract situations induces a relation between *discourse* situations: situation  $z$  is informationally included in situation  $s'$  iff the abstract situation that completely characterizes  $z$  is informationally included in the abstract situation that completely characterizes  $s'$ . I use the relation  $\leq$  to model informational inclusion,  $i, j \leq j'$  indicates that the abstract situation that completely characterizes  $i$  is a subset of the abstract situation that completely characterizes  $j'$ .

Another primitive notion is that of *constituent* in a situation. Informally, an object is a constituent of a situation if it is 'part' of that situation in the real world. A sufficient (but not necessary) condition for an object to be a constituent of a situation is for that object to *live* into the composition of the denotation of some prepositional true at that situation. Below I use the relational symbol  $\vdash$  to indicate constituency.

I use set of temporal relation between situations, including the relations BEFORE (also written  $<$ ), that expresses temporal precedence, and  $\Theta$ , that expresses temporal co-occurrence.

My ontology of situations is in large part derived from that proposed in Epistemic Logic; one extension is the inclusion of *courses of action*.<sup>7</sup> A course of action is an object that behaves like a situation in many ways, yet it actually consists of a group of situations ordered in a sequence and that 'stay together' according to some 'forming principle'. What makes a sequence of events into a course of action may vary—it may be the agent's perception that they form a causal chain, or that some particular individual plays the agent's role in all of them, or some additional factor.<sup>8</sup> Courses of actions in the ontology of CRT are groups of situations. Because the events in a course of action form a sequence, we can define functions  $\text{begin}(e, \text{end})$  and  $\text{next}(e, \text{next})$  which return, for each event  $e$  in a course of action  $\text{coa}$ , the previous event and the next event in the sequence, respectively. Another feature of courses of action that I use below is that for any two successive events  $e_1, e_2$  in a course of action, a function  $\text{B}(e_1, e_2)$  can be defined which returns the time interval between the culminations of the two events.

#### 4.1.4 The Model

I assume that the universe  $U$  is the union of (at least) two disjoint sets, the set  $O$  of objects and the set  $S$  of situations, both of which are actually semi-lattices.<sup>9</sup> The interpretation of types with respect to  $U$  is defined as follows:

- $D_{\text{type } o} = U$ ,
- $D_{\text{type } s} = S$ ,
- $D_{\text{type } i} = \{0, 1\}$ ,
- $D_{\text{type } \tau} = D_{\text{type } o}^n$ .

<sup>7</sup> The term *course of events* was introduced by Barwise and Perry in [1983].

<sup>8</sup> The phenomenon of people forming 'stories' out of descriptions of events, has been studied extensively in the work on understanding narratives [Nelson-Le Gall 1988, Wether 1987, Hoang and Schubert 1992, Karmayev et al. 1993] and it is usually accepted that interventional factors are involved.

<sup>9</sup> A more explicit and analysis of the universe of discourse can be found in [Hoang 1992].

In the rest of this discussion, I generally drop the indication of the universe of the a phrase (e.g., I write  $D_{\text{type } s}$  instead of  $D_{\text{type } s}^U$ ). The model of interpretation for CRT expressions is the following triple:

$$\langle U, S, I \rangle$$

The interpretative function  $I$  assigns an interpretation to constants of type  $\tau$ . As all constants are situation-dependent in CRT, the interpretation of a constant of type  $\tau$  is a function from  $S$  to the domain of  $\tau$ .

$$I: CE_{\text{type } \tau} \rightarrow (S \rightarrow D_{\text{type } \tau}).$$

#### 4.2 DENOTATION AND ENTAILMENT

What gets assigned a value by a model in CRT is a distinguished expression called *Root DRS*, as in DRT. As discussed in Chapter 3, one difference from DRT is that in CRT the Root DRS provides a characterization of a discourse situation as opposed to a characterization of a described situation. Most importantly, the value of a Root DRS  $K$  is a set of functions from situations to truth values. Root DRSs are defined as follows:

**Definition 4.1.** A Root DRS  $K$  is a DRS (i) that contains no unbound variables, (ii) that contains the expression  $\text{[dr]} = s'$  among its conditions, and (iii) such that no DRS and no situation descriptions, assumed in  $K$ , contains that expression among its conditions.

The evaluation function  $\| \cdot \|_{M,d}$  assigns a value to a Root DRS  $K$  with respect to a model  $M$  of evaluation and a discourse situation  $d$  using an auxiliary function  $\| \cdot \|_{M,g,d,c,C}$  that assigns to a meaningful expression  $\alpha$  its value with respect to  $M$ ,  $d$ , a variable assignment  $g$ , a case  $c$ , and a case set  $C$ .

- Let  $K$  be a root DRS,  $\| \cdot \|_{M,d}$  is the set of functions  $\| \cdot \|_{M,g,d,c,C}$ , where  $g$  is an arbitrary assignment,  $c$  is the 'empty' case and defined over any discourse marker, and  $C$  the 'empty' case set - the set of functions that map each situation onto the empty case.

The discourse situation is a situation in  $S$ , used to assign a value to indexical expressions and contextually dependent elements such as pronouns. In particular, a discourse situation must provide values for two indexicals: self and other, that denote the participant to the conversation whose 'believed state' is being modeled and the other participant to the conversation, respectively. The last two parameters of evaluation in  $\| \cdot \|_{M,g,d,c,C}$  have the following interpretation. The case  $c$  is a partial function from discourse markers to situation constituents. As discussed in Chapter 5, discourse markers are used to do the work done by discourse referents in DRT, discourse markers are distinct from variables, whose value is determined by the variable

<sup>10</sup> As discussed in the next section, the term *dr* denotes the discourse situation at which an expression is evaluated, whereas the term *d* denotes the situation of evaluation.

<sup>11</sup> I use the letter  $K$  also printed (K) and with subscripts (K<sub>i</sub>) to indicate DRSs, after [Kamp and Reyle 1993].

assignment  $g$ . This latter is a function from variables of type  $\tau$  to functions from situations to objects of type  $D_\tau$ . The case set  $C$  is a function from situations to cases, used in evaluating discourse markers introduced inside situation descriptions, as discussed below.

The following constraint is imposed on both cases and variable assignments, for all variables  $\alpha$  of type  $\tau$ ,  $g(\alpha)$  is a constant function  $f$ , that is, there is an object  $a \in D_\tau$  such that given any situation  $s$ ,  $f(s)$  is either undefined (if  $a$  is not a constituent of  $s$ ) or else  $f(s) = a$ . All cases as well must be constant functions.

Entailment 'proper' in CRT is a relation between two met DRSS, defined as follows:

**Definition 4.2.** The met DRSS  $K$  entails the met DRSS  $K'$  with respect to model  $M$  and discourse situation  $d$  written  $K \vdash_M d K'$  if for any interpretation  $f \in \llbracket K \rrbracket^M d$  there is an interpretation  $f' \in \llbracket K' \rrbracket^M d$  such that  $f \vdash_M d f'$ .

The notion of entailment with respect to an interpretation  $f \vdash_M d f'$  is defined as follows:

**Definition 4.3.** Let  $f$  and  $f'$  be functions of type  $S \rightarrow \{0, 1\}$ . Then  $f \vdash_M d f'$  iff for every  $s$  such that  $f(s) = 1$ ,  $f'(s) = 1$ .

It is also useful to have a notion of entailment between any two expressions of type  $t$ :

**Definition 4.4.** Let  $\Phi$  and  $\Psi$  be meaningful expressions of type  $t$ . Then  $\Phi \vdash_M d \Psi$  iff for each  $f \in \llbracket \Phi \rrbracket^M d \in C$  there is an  $f' \in \llbracket \Psi \rrbracket^M d \in C$  s.t.  $f \vdash_M d f'$ .

It will be useful to assign a value to expressions of type  $t$  that include 'unbound' discourse markers. I propose the following definition:

Let  $\Phi$  be a CRT expression of type  $t$  in which the discourse markers  $d_1, \dots, d_n$  occur, that are not 'bound' by any DRSS or situation description. (See below.) Then,

$$\llbracket \Phi \rrbracket^M d = \llbracket \frac{d_1 \cdot \dots \cdot d_n}{\Phi} \rrbracket^M d$$

### 4.3 DEFINITION OF THE EVALUATION FUNCTION

Throughout this section I have followed the practice of comparing the clauses specifying the semantics of CRT expressions to the clauses for the analogous expressions found in Drwaty, Wall and Peters [1981], whenever analogous expressions exist. I hope in this way to give the reader a clearer understanding of the intuitions that these recursive truth definitions are meant to capture.

#### 4.3.1 Variables, Non-Logical Constants, Discourse Markers, and Parameters

The values of variables and non-logical constants in Minitague's system are defined as follows:

1. For  $\alpha$  a variable of type  $\tau$ ,  $\llbracket \alpha \rrbracket^M g, w, l = g(\alpha)$ ,  $g(\alpha) \in D_\tau$
2. For  $\alpha$  a constant of type  $\tau$ ,  $\llbracket \alpha \rrbracket^M g, w, l = \text{Fix}(\lambda(w, l)) \in D_\tau$

The denotation of variables and constants stays essentially the same in CRT (modulo the generalizations discussed above), but two new kinds of basic expressions are introduced: discourse markers and parameters. All basic expressions denote sets of functions from situations to their denotation in Intensional Logic. The result is as follows:

1. For  $\alpha$  a variable of type  $\tau$ ,  $\llbracket \alpha \rrbracket^M g, d, c, C =$  the singleton set  $\{g(\alpha)\}$ , where  $g(\alpha)$  is a constant function  $l : S \rightarrow D_\tau$ .
2. For  $\alpha$  a constant of type  $\tau$ ,  $\llbracket \alpha \rrbracket^M g, d, c, C =$  the singleton set  $\{l(\alpha)\}$ , where  $l(\alpha)$  is a function  $f : S \rightarrow D_\tau$ .
3. For  $\alpha$  a discourse marker of type  $\tau$ ,  $\llbracket \alpha \rrbracket^M g, d, c, C =$  the singleton set  $\{c(\alpha)\}$  where  $c(\alpha)$  is a function  $f : S \rightarrow D_\tau$ .
4. For  $\alpha$  a parameter of type  $\tau$ ,  $\llbracket \alpha \rrbracket^M g, d, c, C = \{f \mid f : S \rightarrow D_\tau \text{ is the function such that for each } s, f(s) = a, \text{ where } a \text{ is an object of } \tau, a \in D_\tau, \text{ if } a \dots s, f(s) = \text{undefined otherwise}\}$

The relation between these definitions and the standard ones provided above should be easy to see, but I will give some examples. First of all, the denotation of the constant  $f$  of type  $e$  that may be used to translate the proper name "John"<sup>12</sup> is not an object in  $U$ , but a singleton set, consisting of a unique function  $f$  from situations to objects of type  $e$ , this value is provided by the case  $c$ :

$$\llbracket f \rrbracket^M g, d, c, C = \{f\}, \text{ where } f = l(j)$$

The main difference between CRT and Intensional Logic, as far as non-logical constants are concerned, is that a constant may be undefined at a situation  $s$ , and typically will. A non-logical constant such as CAR denotes in CRT a singleton set of functions from situations to objects of type  $\langle e, e \rangle$ .<sup>13</sup>

$$\llbracket \text{CAR} \rrbracket^M g, d, c, C = \{f\}, \text{ where } f = \lambda(CAR)$$

<sup>12</sup> As discussed in Chapter 5, I translate here proper names as constants to avoid the questions involved in 'anchoring' discourse markers.

<sup>13</sup> Note that relations and properties are already assigned sets of values as their denotation by the evaluation function of CRT, so a treatment of lexical ambiguity could easily be built on top of CRT. I ignore the issue of lexical ambiguity in this discussion.

A bit more interesting is the denotation of a parameter like  $x_T$ , since the power of the system is here put to some use for the first time

$$\|x_T\|_{M,g,d,c} = \{f\} \text{ for an object in } D_x, \text{ and } \|x_T\|_{M,g,d,c} = g' \text{ where } g' \text{ is one of the discourse topics' of } d, f \text{ is the constant function such that } f(s) = a \text{ if } g' \leq s, f(s) = \text{undefined otherwise} \}$$

As discussed in §3.4 parameters are used to represent semantically ambiguous expressions whereas all other terms denote singleton sets, parameters denote non-singleton sets of functions. Each of these functions is constant—it maps all situations onto the same object of the appropriate type. For example, if the subset of  $D_x$  in  $d$  consists of the two atoms  $j$  and  $b$ , then  $\|x_T\|_{M,g,d,c} = \{f, f'\}$ , where  $f$  is the function that maps each situation of which  $j$  is a constituent into  $j$ , and  $f'$  is the function that maps each situation of which  $b$  is a constituent into  $b$ .

There are some distinguished constants in CRT, among these,  $g^*$ , an indexical expression of type  $s$  denoting the current situation of evaluation, and  $ds^*$ , an indexical denoting the discourse situation. Their values are specified as follows.

- $g^*$  is an indexical expression of type  $s$ , and  $\|g^*\|_{M,g,d,c} = \{f\}$ , where  $f$  is the function such that  $f(s) = s$
- $ds^*$  is an indexical expression of type  $s$ , and  $\|ds^*\|_{M,g,d,c} = \{f\}$ , where  $f$  is the function such that  $f(s) = d$
- $self$  is an indexical expression of type  $s$ , and  $\|self\|_{M,g,d,c} = \{f\}$ , where  $f$  is the function such that  $f(s) = self$ , a distinguished element of the discourse situation  $d$
- $other$  is an indexical expression of type  $s$ , and  $\|other\|_{M,g,d,c}$  is the singleton set  $\{f\}$ , where  $f$  is the function such that  $f(s) = \{other\}$ , a distinguished element of the discourse situation  $d$

### 4.3.2 Application

The semantics of application is defined by Dowty, Wall and Peters as follows

- If  $\alpha$  is an expression of type  $\langle a, b \rangle$ , and  $\beta$  an expression of type  $a$ , then  $\alpha(\beta)$  is an expression of type  $b$ , and  $\|\alpha(\beta)\|_{M,g,w,t} = \|\alpha\|_{M,g,w,t}(\|\beta\|_{M,g,w,t})$

The clause specifying the value assigned to the expression  $\alpha(\beta)$  in CRT again differs from the clause just given in that both  $\alpha$  and  $\beta$  denote sets, and also because  $\alpha$  and/or  $\beta$  may be undefined at some situations. The definition that results is as follows

- If  $\alpha$  is an expression of type  $\langle a, b \rangle$ , and  $\beta$  an expression of type  $a$ , then  $\alpha(\beta)$  is an expression of type  $b$ , and  $\|\alpha(\beta)\|_{M,g,d,c} = \{f\}$  for  $\alpha_1 \in \|\alpha\|_{M,g,d,c}$  and  $\beta_1 \in \|\beta\|_{M,g,d,c}$ .

$$f(s) = \begin{cases} \text{undefined} & \text{if either } \alpha_1(s) \text{ or } \beta_1(s) \text{ are undefined;} \\ \{\alpha_1(s)(\beta_1(s))\} & \text{otherwise.} \end{cases}$$

This clause is interesting because application is the first case of an expression whose denotation is obtained by combining the denotations of potentially ambiguous expressions. The technique used here is also used in the other clauses below: the value of an expression like  $\alpha(\beta)$  consists of a set of functions, one function per distinct pair of functions in the denotations of  $\alpha$  and  $\beta$ , the value assigned by the function  $f$  to the situation  $s$  is defined by applying a certain operation (in this case, application) to the values assigned to  $s$  by the functions  $\alpha_1$  and  $\beta_1$ . Thus, if both the denotation of  $\alpha$  and the denotation of  $\beta$  are singleton sets, the denotation of  $\alpha(\beta)$  is also a singleton set, otherwise, ambiguity 'multiplies', as it were.

For example, let us say that  $\|CAR\|_{M,g,d,c} = \{f\}$ , where  $f$  is the function that maps a situation  $s$  in the denotation of the property  $CAR$  in  $g$ , and  $\|j\|_{M,g,d,c} = \{f'\}$ , where  $f'$  is the function that maps each situation  $s$  to an object  $a$  in that situation, both functions may be undefined over some situations. Then  $\|CAR(j)\|_{M,g,d,c}$  is the set  $\{f''\}$ , where  $f''$  is the function such that:  $f''(s)$  is undefined if either  $f(s)$  or  $f'(s)$  are undefined, and  $f''(s) = \{f(s)(f'(s))\}$  otherwise.

More interesting is the case of the expression  $CAR(x_T)$ . Consider again the model used in the example above, where the subset of  $D_x$  in  $d$  is  $\{j, b\}$ . Then,  $\|CAR(x_T)\|_{M,g,d,c} = \{f, f'\}$ , where  $f$  is the function that sends each situation where both  $CAR$  and  $j$  are defined to the value of  $CAR(j)$ , whereas  $f'$  is the function that sends those situations where both  $CAR$  and  $b$  are defined to the value of  $CAR(b)$ .

### 4.3.3 Equality

The denotation of equality statements in Intensional Logic is specified by the following clause

- If both  $\alpha$  and  $\beta$  are expressions of type  $\tau$ ,  $\|\alpha = \beta\|$  is an expression of type  $t$ , and  $\|\alpha = \beta\|_{M,g,w,t} = 1$  iff  $\|\alpha\|_{M,g,w,t} = \|\beta\|_{M,g,w,t}$

This clause is replaced by the following one in CRT:

- If both  $\alpha$  and  $\beta$  are meaningful expressions of type  $\tau$ ,  $\|\alpha = \beta\|$  is a meaningful expression of type  $t$ , and  $\|\alpha = \beta\|_{M,g,d,c} = \{f\}$  for  $h$  in  $\|\alpha\|_{M,g,d,c}$  and  $h'$  in  $\|\beta\|_{M,g,d,c}$ .

$$f(s) = \begin{cases} \text{undefined} & \text{if either of } h(s) \text{ and } h'(s) \text{ is undefined;} \\ 1 & \text{iff } h(s) = h'(s); \\ 0 & \text{otherwise} \end{cases}$$

The most interesting cases of equality statements are those in which one of the elements being equated is a parameter, such as  $\{x_s = b\}$ . Consider again the simple model used above, where  $D_p = \{j, b\}$  and  $\| \cdot \|_{M, g, d, c} = \{f, t\}$ , where  $f$  is the function that maps each situation of which  $x_s$  is a constituent into  $j$ , and  $t$  is the function that maps each situation of which  $b$  is a constituent into  $b$ . Then  $\|x_s = b\|_{M, g, d, c} = \{h, h'\}$ , where  $h$  is the function such that  $h(s) = 0$  if  $f(s)$  is defined, and otherwise  $f(s)$  is undefined, whereas  $h'(s) = 1$  if  $t(s)$  is defined, otherwise  $h'(s)$  is undefined.

This example is important because it shows how adding new statements leads to disambiguation, while the parameter  $x_s$  may refer to any object of its type, in a theory that contains the statement  $\{x_s = b\}$  only one of the possible interpretations of  $x_s$  does not result in a contradiction (a function all of whose values are either 0 or undefined), namely, the interpretation that assigns to  $x_s$  the value of  $b$  in those situations in which  $b$  is defined. Put in other words, that is the only consistent interpretation left for a theory containing such a statement.

At first, this may give the impression that I am now adopting a different definition of semantic ambiguity from the one discussed in Chapter 1. We do need to refine the definition of semantic ambiguity used so far, but the new definition is consistent with the spirit, if not the letter, of the definition proposed in Chapter 1. The idea is that in 'checking' whether a sentence is semantically ambiguous, inconsistent interpretations (i.e. functions that do not assign 1 to any situation in the domain) count only if there are no consistent interpretations available. The revised definition of semantic ambiguity is as follows.

**Definition 4.5** A sentence is *semantically ambiguous* if it has more than one consistent interpretation.

#### 4.3.4 Connectives

The denotation for conjunction provided by Dowty, Wall and Peters is as follows:

- If both  $\alpha$  and  $\beta$  are expressions of type  $t$ ,  $\{\alpha \wedge \beta\}$  is a meaningful expression of type  $t$ , and  $\| \alpha \wedge \beta \|_{M, g, d, c} = \{f\}$  if  $\| \alpha \|_{M, g, d, c} = 1$  and  $\| \beta \|_{M, g, d, c} = 1$ ,  $\| \alpha \wedge \beta \|_{M, g, d, c} = 0$  otherwise.

In (8) we have instead the following clause

- If both  $\alpha$  and  $\beta$  are expressions of type  $t$ ,  $\{\alpha \wedge \beta\}$  is an expression of type  $t$ , and  $\| \alpha \wedge \beta \|_{M, g, d, c} = \{f\}$  for  $h$  in  $\| \alpha \|_{M, g, d, c}$  and  $h'$  in  $\| \beta \|_{M, g, d, c}$ .

$$\| \alpha \wedge \beta \|_{M, g, d, c} = \begin{cases} 0 & \text{if either } h(g) = 0 \text{ or } h'(g) = 0, \\ \text{undefined} & \text{if either } h(s) \text{ or } h'(s) \text{ is undefined} \\ 1 & \text{otherwise} \end{cases}$$

Again, this definition of conjunction leads to a notion of disambiguation as a reduction in the number of non-contradictory interpretations, rather than a process in which interpretations are eliminated altogether. Note that an 'intersective' definition of the semantics of conjunction would not work: we don't want, in other words, a definition of conjunction like the following

- If both  $\alpha$  and  $\beta$  are expressions of type  $t$ ,  $\{\alpha \wedge \beta\}$  is an expression of type  $t$ , and  $\| \alpha \wedge \beta \|_{M, g, d, c} = \| \alpha \|_{M, g, d, c} \cap \| \beta \|_{M, g, d, c}$ .

The problem with this definition is that a conjunction like  $\{\phi \wedge \psi\}$  denotes a non null set only when there is a function in common between the interpretation of  $\phi$  and  $\psi$  which need not be the case in general.

The above definition gives to conjunction in CRT the properties of classical conjunction. So, conjunction in CRT is associative ( $\| \phi \wedge \psi \wedge \chi \|$  is equivalent to  $\| \phi \wedge (\psi \wedge \chi) \|$ ), and conjunction reduction is supported ( $\| \phi \wedge \psi \|$  entails  $\phi$ ) as well as conjunction introduction ( $\phi$  entails  $\| \phi \wedge \psi \|$ ).

Now that we have seen how the clause specifying the denotation of conjunctions, those specifying the denotation of negation and disjunction should be easy to understand.<sup>14</sup>

- If  $\alpha$  is an expression of type  $t$ ,  $\neg \alpha$  is an expression of type  $t$ , and  $\| \neg \alpha \|_{M, g, d, c} = \{f\}$  for  $h$  in  $\| \alpha \|_{M, g, d, c}$ .

$$\| \neg \alpha \|_{M, g, d, c} = \begin{cases} \text{undefined} & \text{if } h(g) \text{ is undefined,} \\ 1 & \text{if } h(g) = 0 \\ 0 & \text{otherwise} \end{cases}$$

}

- If both  $\alpha$  and  $\beta$  are expressions of type  $t$ ,  $\{\alpha \vee \beta\}$  is an expression of type  $t$  and  $\| \alpha \vee \beta \|_{M, g, d, c} = \{f\}$  for  $h$  in  $\| \alpha \|_{M, g, d, c}$  and  $h'$  in  $\| \beta \|_{M, g, d, c}$ .

$$\| \alpha \vee \beta \|_{M, g, d, c} = \begin{cases} 1 & \text{if } h(g) = 1 \text{ or } h'(g) = 1; \\ \text{undefined} & \text{if both } h(g) \text{ and } h'(g) \text{ are undefined, or one is undefined} \\ & \text{and the other is 0,} \\ 0 & \text{otherwise} \end{cases}$$

}

These definitions are fairly standard, but they represent a way of looking at the semantics of connectives that is in contrast with the idea adopted in certain versions of Situation Theory that

<sup>14</sup>The version of negation defined here is the one that Fossard *et al.* (1997) call 'weak' negation. Although something like these 'weak' negations could be defined in CRT, currently only classical negation is used.

information can only be 'extracted' from a situation if that situation explicitly supports it, and therefore a situation that supports the basic facts  $A$  and  $B$  need not support the propositions  $A \wedge B$  or  $A \vee B$ .<sup>11</sup> This property is especially useful when using situation (types) to model attitude reports, in that it avoids the logical omniscience problem (for some discussion of this problem, see [Levesque, 1990] and [Eggen and Halpern, 1988]). CRT suffers from the logical omniscience problem to some extent.<sup>12</sup>

#### 4.3.5 Lambda Abstraction

The clause specifying the denotation of lambda-abstraction in Dowty, Wall and Peters is the following:

- If  $\alpha$  is a variable of type  $\tau$  and  $\beta$  a meaningful expression of type  $\tau'$ , then  $\lambda \alpha \beta$  is an expression of type  $(\tau, \tau')$ , and  $\|\lambda \alpha \beta\|_{M,g,w}$  is that function  $h: D_\tau \rightarrow D_{\tau'}$  such that for all objects  $a$  in  $D_\tau$ ,  $h(a)$  is equal to  $\|\beta\|_{M,g(n/a),w}$ .

Some care is required in order to get a semantics for lambda-abstraction in CRT that preserves properties such as  $\beta$ - and  $\eta$ -reduction. The following clause, for example, does not do the job:

- If  $\alpha$  is a variable of type  $\tau$  and  $\beta$  a meaningful expression of type  $\tau'$ , then  $\lambda \alpha \beta$  is an expression of type  $(\tau, \tau')$ , and  $\|\lambda \alpha \beta\|_{M,g,d,c,C}$  is the set  $\{f \mid \text{where } f: S \rightarrow (D_\tau \rightarrow D_{\tau'}) \text{ and for all situations } s, f(s) = h: D_\tau \rightarrow D_{\tau'}, \text{ such that, for all objects } a \text{ in } D_\tau, h(a) \text{ is equal to } h'(s), \text{ where } h' \in \|\beta\|_{M,g(n/a),d,c,C}\}$ .

It is easy to show that lambda-abstraction defined in this way does not preserve  $\eta$ -reduction. Imagine for example that  $S = \{s_1, s_2\}$ ,  $D_\tau = \{a, b\}$ , and the expression  $\beta$  of type  $(\tau, \tau')$  has the following denotation

$$\left\{ \left\{ \begin{array}{l} s_1 \rightarrow \left[ \begin{array}{l} a \rightarrow \frac{a_1}{f_1} \\ b \rightarrow \frac{a_2}{f_1} \end{array} \right], \\ s_2 \rightarrow \left[ \begin{array}{l} a \rightarrow \frac{a_2}{f_2} \\ b \rightarrow \frac{a_1}{f_2} \end{array} \right] \end{array} \right\} \right\}$$

Then  $\|\lambda \alpha \beta\|_{M,g(n/a),d,c,C}$  is as follows

$$\left\{ \left\{ \begin{array}{l} s_1 \rightarrow \frac{a_1}{f_1} \\ s_2 \rightarrow \frac{a_2}{f_2} \end{array} \right\}, \left\{ \begin{array}{l} s_1 \rightarrow \frac{a_2}{f_2} \\ s_2 \rightarrow \frac{a_1}{f_1} \end{array} \right\} \right\}$$

and  $\|\beta\|_{M,g(n/b),d,c,C}$  is as follows.

<sup>11</sup> Situation Theory alleviates in part the logical omniscience problem, in that logical propositions such as  $(2 \times 2 = 4)$  need not be defined over all situations, and therefore the consequences of a proposition  $p$  that is true at a situation  $s$  with a boundary need not be defined over  $s$ . However, if  $p$  is false at  $s$ ,  $s$  will support the consequences of  $p$  with all logical contradictions.

$$\left\{ \left\{ \begin{array}{l} s_1 \rightarrow \beta_1 \\ s_2 \rightarrow \beta_1 \end{array} \right\}, \left\{ \begin{array}{l} s_1 \rightarrow \beta_2 \\ s_2 \rightarrow \beta_2 \end{array} \right\} \right\}$$

Then, under the definition above, the denotation of  $\lambda \alpha \beta(\alpha)$  will contain the following function, that is not part of the denotation of  $\beta$  (hence,  $\eta$ -reduction is not a sound inference rule)

$$\left\{ \begin{array}{l} s_1 \rightarrow \left[ \begin{array}{l} a \rightarrow \frac{a_1}{f_1} \\ b \rightarrow \frac{a_2}{f_2} \end{array} \right] \\ s_2 \rightarrow \left[ \begin{array}{l} a \rightarrow \frac{a_2}{f_2} \\ b \rightarrow \frac{a_1}{f_1} \end{array} \right] \end{array} \right\}$$

Intuitively, the problem with the definition above is that it does not 'preserve' the functions in the denotation of  $\beta$ . A definition of lambda-abstraction that does preserve these functions, and therefore preserves the soundness of  $\beta$ - and  $\eta$ -reduction, can be obtained as follows. First of all, a class of functions  $\Delta_\alpha$  is defined, each of which maps a CRT expression  $\psi$  of type  $\tau'$ , in which the variable  $\alpha$  of type  $\tau$  may occur free, and such that  $\|\psi\|_{M,g,d,c,C} = \{f_1, \dots, f_{n_\alpha}\}$ , into the set of functions  $U = \{u_1, \dots, u_{n_\alpha}, \dots\}$ ,  $u_i: (D_\tau \rightarrow (S \rightarrow D_{\tau'}))$ . There is one  $u_i$  in  $U$  for each  $f_i$  in the denotation of  $\psi$ ; the function  $u_i$  'keeps track' of the value that would be assigned to  $\psi$  relative to  $M$  and  $d$ , depending on the value of  $f_i(\alpha)$ , according to the interpretation of  $\psi$  represented by  $f_i$ . The function  $\Delta_\alpha$  can then be used to define the semantics of  $\lambda \alpha \beta$  as follows:

- Let  $\alpha$  be a variable of type  $\tau$ , and  $\beta$  a meaningful expression of type  $\tau'$ . Then  $\lambda \alpha \beta$  is a meaningful expression of type  $(\tau, \tau')$ , and  $\|\lambda \alpha \beta\|_{M,g,d,c,C} = \{f \mid f = \omega(h), \text{ for } h \in \Delta_\alpha(\beta)^{M,g,d,c,C}\}$ .

The function  $\omega$  used in the clause above maps a function  $h: (D_\tau \rightarrow (S \rightarrow D_{\tau'}))$  into a function  $f: (S \rightarrow (D_\tau \rightarrow D_{\tau'}))$ . The definition of  $\omega$  is also given below.

The clauses that provide a definition for  $\Delta_\alpha$  that ensures that the semantics of lambda-abstraction has the desired properties can be derived from those that give the definition of  $\|\cdot\|$ . Once the main idea is understood, therefore I only specify some of the most important clauses here, including those specifying the value of  $\Delta_\alpha$  for application and lambda-abstraction. The base of the recursion is provided by the following clauses, specifying the value of  $\Delta_\alpha$  for  $\alpha$  and for expressions that do not contain  $\alpha$ :

- $\Delta_\alpha(\alpha)^{M,g,d,c,C} = \{f\}$ , where  $f$  is the function such that, for every  $a \in D_\tau$ ,  $f(a) = h_a$ , where  $h_a$  is the constant function such that for every  $s$ ,  $h_a(s) = a$  if  $a \in s$ , undefined otherwise.
- If  $\beta$  is of type  $\tau'$ , and  $\alpha$  does not occur in  $\beta$ ,  $\Delta_\alpha(\beta)^{M,g,d,c,C} = \{f \mid \text{given a function } h \in \|\beta\|_{M,g,d,c,C}, h: S \rightarrow D_{\tau'}, f \text{ is the constant function } (D_\tau \rightarrow (S \rightarrow D_{\tau'})) \text{ such that for all } a \in D_\tau, f(a) = h\}$ .

$\Delta_\alpha(\alpha)^{M,g,d,c,C}$  is a function that, given a possible value for  $\alpha$ , returns a function that maps each situation on that value. If  $\beta$  is an expression that does not contain  $\alpha$ ,  $\|\beta\|_{M,g,d,c,C}$  does not depend on the value of  $\alpha$ , and therefore  $\Delta_\alpha(\beta)^{M,g,d,c,C}$  is the set of interpretations that  $\beta$  may have. For example, say that  $\alpha$  is of type  $e$ , and  $\beta$  is of type  $(e, t)$ . Let us ignore the structure of the domain, and just say that  $D_e = \{a, b, s, s', s''\}$ . Say that  $a \in s$  and  $s' \in s$ , and  $b \in s'$  and  $s'' \in s$ . Let us also assume that  $\beta$  is ambiguous, as follows

$$\bullet \parallel \lambda \parallel M.g.d.c.C = \{h = \left[ \begin{array}{c|c} \left[ \begin{array}{c} a \rightarrow 1 \\ b \rightarrow 0 \\ c \rightarrow 0 \end{array} \right] & \left[ \begin{array}{c} s \rightarrow \\ s' \rightarrow \end{array} \right] \\ \hline \left[ \begin{array}{c} a \rightarrow 1 \\ b \rightarrow 1 \\ c \rightarrow 0 \end{array} \right] & \left[ \begin{array}{c} s \rightarrow \\ s' \rightarrow \end{array} \right] \end{array} \right\}$$

We get the following values of  $\Delta_\alpha$  from the definitions above

$$\bullet \Delta_\alpha(\alpha)^{M.g.d.c.C} = \left\{ \left[ \begin{array}{c} s \rightarrow a \\ s' \rightarrow \text{undefined} \\ s'' \rightarrow a \\ s''' \rightarrow b \\ s'''' \rightarrow b \end{array} \right] \right\}$$

$$\bullet \Delta_\alpha(\alpha)^{M.g.d.C} = \left\{ \left[ \begin{array}{c} a \rightarrow b \\ b \rightarrow b \end{array} \right] \right\}$$

Next, let us consider application

- If  $f$  is of type  $\tau''$ ,  $\tau'$ , and  $h$  is of type  $\tau''$ ,  $\Delta_\alpha((f(h))^{M.g.d.C}) = \{f \mid f \text{ is a function s.t. for all } a \in D_{\tau'}, f(a) = \text{the function } h \text{ s.t. for some } b \in D_{\tau'} (S \rightarrow (D_{\tau'} \rightarrow D_{\tau'}) \cap \Delta_\alpha)^{M.g.d.C} \text{ and } b' \in D_{\tau'} (S \rightarrow D_{\tau'}) \cap \Delta_\alpha^{M.g.d.C}, h(s) \text{ is undefined if any of } f(a) [f'(a)](s), f''(a) \text{ and } [f''(a)](s) \text{ is undefined, else } h(s) = [f'(a)](s) [f''(a)](s) \}$

With these definitions, and assuming the values for  $\Delta_\alpha(\alpha)$  and  $\Delta_\alpha(\alpha)$  above, we get the following value for  $\Delta_\alpha(f(h))$

$$\bullet \Delta_\alpha(f(h))^{M.g.d.C} = \left\{ \left[ \begin{array}{c} a \rightarrow \left[ \begin{array}{c} s \rightarrow 1 \\ s' \rightarrow 1 \end{array} \right] \\ b \rightarrow \left[ \begin{array}{c} s \rightarrow 0 \\ s' \rightarrow 1 \end{array} \right] \end{array} \right] \right\}$$

It can be seen that the definition of  $\Delta_\alpha$  closely mimics the definition of  $\parallel \parallel$ ; this guarantees that the set of functions in  $\Delta_\alpha(\Psi)^{M.g.d.C}$  is going to be isomorphic to the set of functions in  $\parallel \Psi \parallel^{M.g.d.C}$ . Another example of this parallelism is provided by the clause specifying the value of  $\Delta_\alpha$  in the case of a (type  $t$ ) conjunction, to be compared with the clause specifying the value of  $\parallel \parallel$  for a conjunction.<sup>16</sup>

<sup>16</sup> This clause can easily be modified to get a definition of generalized conjunction.

- If  $\Psi$  and  $\Phi$  are of type  $t$ ,  $\Delta_\alpha(\Psi \wedge \Phi)^{M.g.d.C} = \{f \mid f \text{ is a function s.t. for all } a \in D_{\tau'}, f(a) = \text{the function } h \text{ s.t. } h(s) = 0 \text{ if either } [f'(a)](s) = 0 \text{ or } [f''(a)](s) = 0, \text{ else } h(s) \text{ is undefined if either } [f'(a)](s) \text{ or } [f''(a)](s) \text{ is undefined, else } h(s) = 1, \text{ for } b' \in D_{\tau'} \rightarrow (S \rightarrow \{0,1\}) \in \Delta_\alpha^{M.g.d.C} \text{ and } b'' \in D_{\tau'} \rightarrow (S \rightarrow \{0,1\}) \in \Delta_\alpha^{M.g.d.C} \}$

Finally, we have to consider what happens when  $\alpha$  gets bound in a lambda abstraction.<sup>17</sup> Abstracting over  $\alpha$  makes the value of an expression independent from the value of  $\alpha$ . The value of  $\Delta_\alpha$  in this case is specified by the following clause:

- Let  $\alpha$  be a variable of type  $\tau$ , and  $\beta$  a meaningful expression of type  $\tau'$ . Then  $\Delta_\alpha(\lambda \alpha. \beta)^{M.g.d.C} = \{f \mid \text{for some } h \in \parallel \lambda \alpha. \beta \parallel^{M.g.d.C}, f \text{ is the constant function } f(a) = h \}$

The last step is to provide a definition of the function  $\omega$  that turns a function  $h : D_{\tau'} \rightarrow (S \rightarrow D_{\tau'})$  into a function  $f : (S \rightarrow D_{\tau'}) \rightarrow D_{\tau'}$

$$\omega(h) (s_1) = \text{the function } f : (D_{\tau'} \rightarrow D_{\tau'}) \text{ s.t. } f(a_2) = [h(a_2)](s_1)$$

It can be shown that with the definitions of  $\Delta_\alpha$  and  $\omega$  just given, the definition of  $\parallel \lambda \alpha. \beta \parallel^{M.g.d.C}$  coincides with both  $\parallel$  and  $\eta$ -reduction in sound

#### 4.3.6 DRSs

In order to help understanding the semantics of constructs that are not found in Intensional Logic, like DRSs and situation descriptions, in the next few sections I will proceed as follows: I will try, first of all, to convey an understanding of the interpretation of these constructs by leaving aside the issue of ambiguity and specifying the semantics of these constructs by means of clauses that refer to  $\parallel \parallel^{M.g.d.C}$  as if its value were a single function from situations to objects in the domain. I will then give the actual definition.

Some terminology first. The marker set of a DRS  $K$ , written  $MS_K$ , is the set of discourse markers in  $K$ , together with all the discourse markers accessible from  $K$ . The condition set of  $K$ , written  $COND_K$ , is the set of all conditions in  $K$ . A case set  $C$  is said to extend the case set  $C'$  iff every  $h$  s.t.  $h' = C'(s)$ , for some  $s \in S$ , is defined over all the markers on which  $C(s)$  is defined, but may additionally be defined over some discourse markers.

If  $\parallel \parallel^{M.g.d.C}$  were a single-value function, the semantics of DRSs could be specified as in the following clause

- If  $m_1, \dots, m_n$  are discourse markers, and  $\Phi_1, \dots, \Phi_n$  are expressions of type  $t$  then

$$K = \left[ \begin{array}{c} m_1, \dots, m_n \\ \hline \Phi_1 \\ \hline \Phi_n \end{array} \right]$$

<sup>17</sup> As we will see below, binding by quantifiers is defined in terms of binding by lambda-abstraction, therefore this is the only case of binding we have to consider.

is an expression of type  $t$ , and  $\|K\|M.g.d.c.C =$  the function  $f(s)$

$$f(s) = \begin{cases} 1 & \text{if } c \text{ is not defined over } m_1, \dots, m_n, \text{ and there is a case } c' \text{ that takes} \\ & \text{values over } s \text{ and extends } c \text{ by additionally being defined on} \\ & \text{exactly } m_1, \dots, m_n, \text{ and there is a case set } C' \text{ that extends } C, \text{ so} \\ & \text{that for all } \phi_i \text{ in } (CON_k, \|\phi_i\|M.g.d.c.C(s) = 1) \\ 0 & \text{if } c \text{ is defined over one of } m_1, \dots, m_n, \text{ or for all cases } c' \text{ defined over} \\ & \text{MS}_k \text{ and taking values over } s, \text{ and all case sets } C' \text{ extending } C, \\ & \|\phi_i\|M.g.d.c.C'(s) = 0 \text{ for some } \phi_i \text{ in } CON_k \\ & \text{undefined otherwise} \end{cases}$$

This clause says that the expressions in the condition set of a DRS may be thought of as a conjunction, and the main points for using a DRS is that the discourse markers mechanism provides a way to specify that certain referents are new. This latter feature is crucial for the representation of indefinites, for example, the sentence 'A dog came in' gets translated as the following DRS  $K$  (ignoring tense)

$$(4.1) \quad K = \frac{\begin{array}{|c|} \hline m \\ \hline \text{DOG}(m) \\ \hline \text{CAME-IN}(m) \\ \hline \end{array}}{\quad}$$

This representation guarantees that  $m$  is novel.  $\|K\|M.g.d.c.C$ , according to the clause above, is the function  $f(s)$   $f(s) = 1$  if  $m$  is a new discourse marker, and there is a case  $c'$  that is defined over  $MS_k$ , assigning to  $m$  (and to all discourse markers accessible from  $K$ ) values in  $s$ , and a new case set  $C'$ , such that both  $\|\text{DOG}(m)\|M.g.d.c.C(s) = 1$  and  $\|\text{CAME-IN}(m)\|M.g.d.c.C'(s) = 1$   $f(s) = 0$  if  $c$  is defined over  $m$ , or if every extension of  $c$  that includes a value for  $m$  assigns to  $\text{DOX}(m)$  or  $\text{CAME-IN}(m)$  functions whose value in  $s$  is 0

When trying to generalize the clause above to allow for conditions of a DRS to have more than one interpretation, we are faced with a problem similar to the one we had with lambda-abstraction: unless we are careful, we may not 'propagate the ambiguity' of the conditions correctly. The solution is similar to that proposed for lambda-abstraction. First of all, we define a class of auxiliary functions  $\Sigma m_1, \dots, m_n, p_1, \dots, p_n$  that play the same role of the function  $\Delta$ , used to define the semantics of lambda-abstraction  $\Sigma m_1, \dots, m_n, p_1, \dots, p_n, (\phi)$   $\Delta$  returns the set of functions from  $(n+s)$ -tuples of objects of the type of  $m_1, \dots, m_n, p_1, \dots, p_n$  to objects of type  $S \rightarrow D$  ( $t$  being the type of  $\phi$ ) that we obtain by 'abstracting' each of the interpretations of  $\phi$  over the values that might be assigned to  $m_1, \dots, m_n, p_1, \dots, p_n, \Sigma m_1, \dots, m_n, p_1, \dots, p_n$  does for markers, in other words, what  $\Delta$  does for the variable  $n$ —it 'assembles' all of the alternative interpretations of  $\phi$  much as  $\|M.g.d.c.C\|$  would, except that instead of a set of interpretations depending on the value of  $m_1, \dots, m_n, p_1, \dots, p_n$ , we get a set of functions each of which returns, given a tuple of values for  $m_1, \dots, m_n, p_1, \dots, p_n$ , the function that would be returned by  $\|M.g.d.c.C\|$  were  $c$

and  $C$  to assign those values to the markers.<sup>18</sup>

Once we have  $\Sigma m_1, \dots, m_n, p_1, \dots, p_n$ , we can isolate, for each condition occurring in a DRS, one of the interpretations that that condition would have independently from the choice of a case set, and for each permutation of the interpretations of a DRS's conditions we can then obtain one interpretation for that DRS, as follows:

- If  $m_1, \dots, m_n$  are discourse markers, and  $\phi_1, \dots, \phi_n$  are expressions of type  $t$ , then

$$K = \frac{\begin{array}{|c|} \hline m_1, \dots, m_n \\ \hline \phi_1 \\ \hline \phi_n \\ \hline \end{array}}{\quad}$$

is an expression of type  $t$ , and  $\|K\|M.g.d.c.C =$

{f | f given a tuple  $h_1 \in \Sigma m_1, \dots, m_n, p_1, \dots, p_n, (\phi_1), \dots, h_n \in \Sigma m_1, \dots, m_n, p_1, \dots, p_n, (\phi_n)$ , where  $p_1, \dots, p_n$  are all the discourse markers introduced in a situation description  $s, K'$  that is embedded in  $K$  but not in any other DRS inside  $K$ ,

$$f(s) = \begin{cases} 1 & \text{if } c \text{ is not defined over } m_1, \dots, m_n, \text{ and there is a case } c' \text{ taking values} \\ & \text{over } s \text{ and extending } c \text{ by being defined over } m_1, \dots, m_n \text{ (exactly),} \\ & \text{and a case set } C' \text{ extending } C \text{ such that, for all } h_i \text{ participating in} \\ & \text{the definition of } f, h_i((c'(m_1), \dots, c'(m_n), [C(\|s_p, \|M.g.d.c.C')\|Xs) = 1, \text{ where } s_p \text{ is} \\ & (p_1), \dots, [C(\|s_p, \|M.g.d.c.C')\|Xs) = 1, \text{ where } s_p \text{ is} \\ & \text{the term of type } s \text{ such that } p_i \text{ occurs inside } K \text{ in a situation} \\ & \text{description of the form } s_p, K' \\ 0 & \text{if } c \text{ is defined over one of } m_1, \dots, m_n, \text{ or for all cases } c' \text{ taking} \\ & \text{values over } s \text{ and defined over } MS_k, \text{ and case sets } C' \text{ as above,} \\ & h_i((c'(m_1), \dots, c'(m_n), [C(\|s_p, \|M.g.d.c.C')\|Xs) = 1, \dots, [C(\|s_p, \\ & \|M.g.d.c.C')\|Xs) = 0 \text{ for some } h_i,} \\ & \text{undefined otherwise} \end{cases}$$

The definition for  $\Sigma m_1, \dots, m_n, p_1, \dots, p_n$  can be obtained by adapting that for  $\Delta$ , it is in fact simpler, as we don't need to worry about discourse markers being abstracted over.<sup>19</sup>

<sup>18</sup> Note that neither variables nor discourse markers introduce ambiguity, i.e. the definitions above

<sup>19</sup> The reader should be warned that a semantics of DRSs defined in terms of situation extensions is not necessarily equivalent to one defined in terms of embedding extensions. The former, however, example from Heim (1982, 1990) illustrates the fact that semantics defined on the basis of assignments are finer-grained than semantics defined in terms of situations.

(4.2) If a lambda-term meets another name, let it be a function.


### 4.3.7 Situation Descriptions

The main difference between 'simple' DRSs and situation descriptions is that the discourse markers that are 'accessible' from within a DRSs are as specified in [Kamp and Reyle 1993], whereas from within the situation descriptions  $s$   $K$  all discourse markers are accessible that are introduced inside situation descriptions that characterize the situation  $s$  or one of its substitutions. This accessibility between situation descriptions is achieved by using the case-set  $C$  to assign a value to the discourse markers inside a situation description.

Again, I will at first ignore the fact that some of the conditions in the 'DRS' part  $K$  of a situation description might be semantically ambiguous, and I will give a simpler version of the semantics of situation descriptions in order to explain how situation descriptions accomplish their twofold semantic job: to describe the truth at a situation and to allow for more complex forms of accessibility. I will also introduce first a more complex form of situation description written  $\tau_p K$ .

If we were to define  $\| \cdot \|_{M.g.d.c.C}$  as a simple-valued function, the semantics of the situation description  $\tau_p K$  could be given by the following clause

- If  $s$  is a term of type  $s$ ,  $m_1$ ,  $m_2$  are discourse markers, and  $\phi_1$ ,  $\phi_2$  are meaningful expressions of type  $t$ , then




is an expression of type  $t$  and the value of this expression is the function  $f$  such that

The key observation is that in a model in which a belief needs another belief, and in which only one belief believes the other the semantics defined in situation descriptions will predict that (4.2) is true (because only one situation is created) whereas the semantics defined on assignments will predict that (4.2) is false (because two distinct assignments are involved): this latter result is more in agreement with the intuition. I believe that the solution to this problem is to make the evaluation of situations depend on the roles of their constituents, i.e. to treat situations not as representations of the world out there, but as ways of categorizing that world.

$$f(s) = \begin{cases} 1 & \text{if for all } m_1, \dots, m_n, C(\|s\|_{M.g.d.c.C}(s)) \text{ is not defined over } m_j, \text{ whereas the case } c' = C(\|s\|_{M.g.d.c.C}(s)) \text{ extends } C(\|s\|_{M.g.d.c.C}(s)) \text{ by being defined over } m_1, \dots, m_n \text{ (exactly) and assigns them values in } \|c\|_{M.g.d.c.C}(s) \text{ such that that for all } \phi_j \text{ in } CON_{\phi_j} \\ & \| \phi_j \|_{M.g.d.c.C}(s) = 1 \\ & \| \phi_j \|_{M.g.d.c.C}(s) \text{ is defined over one of } m_1, \dots, m_n, \text{ or if } C(\|s\|_{M.g.d.c.C}(s)) \text{ assigns to } m_1, \dots, m_n \text{ values such that for at least one } \phi_j \text{ in } CON_{\phi_j} \\ & \| \phi_j \|_{M.g.d.c.C}(s) = 0, \\ & \text{undefined. If there is one } \phi_j \text{ in } CON_{\phi_j} \text{ s.t.} \\ & \| \phi_j \|_{M.g.d.c.C}(s) \text{ is undefined} \end{cases}$$

The key aspects of this definition are (i) that each alternative interpretation of a statement of the form  $s_p K$  evaluates to 1 in a situation  $s$  iff the value of  $s$  in  $s$  is 'of the type' specified by  $K$ , in the sense that each condition in  $K$  assigns 1 to  $\| \phi \|_{M.g.d.c.C}(s)$  under that interpretation, and (ii) that the case set  $C$  is used to allow for inter-situational accessibility by letting the case used to evaluate the discourse markers in  $\|s\|_{M.g.d.c.C}$  be an expansion of the case used to evaluate the discourse markers in  $\|s'\|_{M.g.d.c.C}$ .

Since the discourse markers introduced in a situation description get their value from  $C$ , all discourse markers that are part of the characterization of  $s'$  are 'accessible' from within  $s$  if  $s_p K$  is the case. For example, consider the DRS in (4.3) containing two situation descriptions. The discourse marker  $m_1$  introduced in the first situation description, can be referred to in the second one, because this latter provides a characterization of the situation  $s'$  that expands  $s$  and therefore, because of the condition on case sets just introduced, the value of  $m_1$  'within the description of  $s'$ ' must be the same as the value it gets 'within the description of  $s$ '.




(4.3)

(Note that 'expansion' in the sense used here is a weaker notion than situation inclusion:  $s \subseteq s'$  does not entail  $s \subseteq s'$ . In fact, the definitions of 'situation descriptions' and 'situation inclusion' do not ensure that entailment proceeds in the other direction, either.)

The property that situation descriptions have of making a discourse marker introduced in the situation description  $s$   $K$  accessible from within the situation description  $s' K'$  that extends



$\tau'$  is particularly important considered that, as discussed in §3.7 and §3.8, the common ground contains a distinct situation description for each utterance, corresponding to the conversational event introduced by that utterance. Thus, a text like "A dog came in. It sat under the table" would result in a common ground of the form



What makes the value of the marker  $m$  in  $s'$  identical to the value in  $s$  is the fact that both the same case is used to assign a value to markers.


Discourse markers introduced 'inside' situation descriptions get existential force, because of the existential quantification over case sets in the clause defining the semantics of DRSs. In addition, the formulation of the semantics of situation descriptions given above, by requiring that  $C(\|s\| M.g.d.c.C(s))$  be defined only on the union of the markers in  $MS_K$ , together with the markers on which  $C(\|s\| M.g.d.c.C(s))$  is defined, prevents the possibility that a CRT theory might include two situation descriptions of the form  $s_j K$  and  $s_j' K'$ , thus imposing a 'novelty condition' on discourse markers first introduced in situation descriptions analogous to that on discourse markers first introduced in DRSs.

The form of situation description most commonly used in the dissertation,  $\tau' K$ , in which the second situational argument is omitted, is in fact a form of shorthand. Assume we have a theory containing the following 'complex' situation descriptions




Since the case for  $s''$  in a case set  $C$  must assign to all of the discourse markers introduced in  $K^1$ ,  $K^{n-1}$  the same value that is assigned to them by the cases for  $s^1, s^2, \dots, s^{n-1}$ , and since each

fact supported by (the denotations of)  $s^1, s^2, \dots, s^{n-1}$  must also be supported by  $\|s''\| M.g.d.c.C$ , we can write the DRS above in the following simplified form:



We can now get to the 'ambiguous' version of the definition of situation description. This is defined by the following clause:

- If  $s$  is a term of type  $s$ ,  $m_1, \dots, m_n$  are discourse markers, and  $\phi_1, \dots, \phi_m$  are meaningful expressions of type  $\tau$ , the




is an expression of type  $\tau$ , and the value of this expression according to  $M.g.d.c$  and  $C$  is the following set:

- (f) given any  $n$  functions  $b_1 \in \|\phi_1\| M.g.d.c.C, \dots, b_m \in \|\phi_m\| M.g.d.c.C$ ,

$$f(s) = \begin{cases} 1 & \text{if } C(\|s\| M.g.d.c.C(s)) \text{ is not defined over } m_1 \cdot m_n, \text{ whereas the} \\ & \text{case } c' = C(\|s\| M.g.d.c.C(s)) \text{ is defined over them (exactly) and} \\ & \text{assigns them values in } \|\phi_1\| M.g.d.c.C(s) \text{ such that for all } \phi_j \text{ in} \\ & \text{CON}_K, \|\phi_j\| M.g.d.c.C(\|s\| M.g.d.c.C(s)) = 1. \\ 0 & \text{if } C(\|s\| M.g.d.c.C(s)) \text{ is defined over one of } m_1 \cdot m_n, \text{ or if} \\ & \text{at least one } \phi_j \text{ in CON}_K \\ & \|\phi_j\| M.g.d.c.C(\|s\| M.g.d.c.C(s)) = 0 \\ \text{undefined} & \text{if there is one } \phi_j \text{ in CON}_K \text{ s.t.} \\ & \|\phi_j\| M.g.d.c.C(\|s\| M.g.d.c.C(s)) \text{ is undefined} \end{cases}$$

Situation descriptions 'partition' the set of facts in a DRS if a situation description  $s K$  contains a condition  $\phi$ , and if that situation description is evaluated with respect to the situation  $s$ ,  $\phi$  is

evaluated with respect to  $\| \lambda \|_{M.g.d.c.C.}$ , not  $g$ . Because of this, situation descriptions can also be used to represent presuppositions, i.e., statements that need to be assumed in order to be able to assign a value to a certain expression, as discussed in §3.6. I discussed there, for example, the fact that the presuppositional aspects of the sentence "The king of France is bald" can be represented by using parameters together with situation descriptions as shown in (4.5).




We have here two situation descriptions, one of which (the one providing a partial characterization of  $x$ ) is embedded inside the other. The expression  $K.O.F(x)$  is not evaluated with respect to the situation  $\| \lambda \|_{M.g.d.c.C.}$ , but with respect to whatever situation is chosen as the interpretation of  $x$ . Furthermore, according to the condition on parameters formulated in §3.4, the 'Condition on Discourse Interpretation', a discourse situation is only felicitous if parametric statements such as (4.5) have been resolved, either by accommodation or by identifying the fact in question. By using situation descriptions together with parameters, we get a theory of presupposition and presupposition cancellation very similar to the one proposed by Heim and van der Sandt.

Event descriptions introduced in §3.5.2 can be modeled as situation descriptions, and their additional 'completeness' properties can be specified by means of meaning postulates as seen in that section. Finally, we can use situation descriptions of the form  $s \models K$  to define the expression  $f(s) \models \Phi$  as an abbreviation

**Definition 4.6**  $f(s) \models \Phi \equiv \text{def } s \models \Phi$

### 4.3.8 Quantifiers

The treatment of quantifiers in CRT is based on Generalized Quantifiers Theory (Barwise and Cooper, 1981; Gärdenfors, 1987; van Eijck, 1991), which is the idea that determiners denote relations between two sets. In CRT the theory of generalised quantifiers is implemented by means of abstractions over DRSs. For example, the translation of the sentence 'Every kid climbed a tree' discussed in Chapter 5 is required below



I will only present here the semantics for the determiner every, the other strong determiners can be similarly defined. If we ignore the fact that  $\| \lambda \|_{M.g.d.c.C.}$  is set valued, the semantics of the construct  $\text{every}(\lambda \alpha \text{ } f) \models K_1 / \lambda \alpha \text{ } K_2$  could be defined in first approximation, as in the following clause

- Let  $K_1$  and  $K_2$  be DRSs, let  $s$  be a term of type  $s$ , and let  $\alpha$  be a variable of type  $e$ . Then  $\text{every}(\lambda \alpha \text{ } f) \models K_1 / \lambda \alpha \text{ } K_2$  is a meaningful expression of type  $t$  and  $\| \text{every}(\lambda \alpha \text{ } f) \models K_1 / \lambda \alpha \text{ } K_2 \|_{M.g.d.c.C.}$  is the function  $f(s)$


$$f(s) = \begin{cases} \text{undefined} & \text{iff either } \lambda \alpha \text{ } f \models K_1 / \lambda \alpha \text{ } K_2 \text{ is undefined or } \lambda \alpha \text{ } f \models K_2 \text{ is undefined} \\ 1 & \text{iff for every } a \text{ s.t. } \lambda \alpha \text{ } f \models K_1 / \lambda \alpha \text{ } K_2 \text{ (s)}(a) = 1 \text{ } \lambda \alpha \text{ } f \models K_2 \text{ (s)}(a) = 1 \\ 0 & \text{otherwise} \end{cases}$$

This definition generalizes well to the case in which we take  $\| \lambda \|_{M.g.d.c.C.}$  to be a set

- $\| \text{every}(\lambda \alpha \text{ } f) \models K_1 / \lambda \alpha \text{ } K_2 \|_{M.g.d.c.C.} = \{ s \mid \text{for } h : \lambda \alpha \text{ } f \models K_1 / \lambda \alpha \text{ } K_2 \text{ and } h' \in \lambda \alpha \text{ } K_2 \text{ } \lambda \alpha \text{ } K_2 \text{ } \lambda \alpha \text{ } K_2 \}$

$$f(s) = \begin{cases} \text{undefined} & \text{iff either } h(s) \text{ is undefined or } h'(s) \text{ is undefined} \\ 1 & \text{iff for every } a \text{ s.t. } h(s)(a) = 1 \text{ } h'(s)(a) = 1 \\ 0 & \text{otherwise} \end{cases}$$

However, the definition above need to be fixed to allow for the kind of 'donkey' anaphora exemplified by sentences such as "Every farmer who owns a donkey beats it", whose CRT translation (ignoring discourse aspects) is shown in (4.7).



What the example shows is that a discourse marker introduced in the restriction of a quantifier may be anaphorically referred to in the nuclear scope, however, with the semantics provided above, neither the discourse marker  $s'$  nor the discourse marker  $\lambda$  are guaranteed to be assigned by the case of  $s'$  the same value they are assigned by the case for  $s$ .

This kind of reference can be accommodated if we 'build' conservativity into the semantics of quantified expressions. Conservativity is a property of relations between sets.

**Definition 4.7** A relation  $R$  is said to be *conservative* iff  $R(A, B) \equiv R(A, A \cap B)$ .

It is known that conservativity is a property of all determiners. If we 'build' conservativity into the definition of every (non ambiguous version), we get the following

- Let  $K_1$  and  $K_2$  be DRSs, let  $s$  be a term of type  $s$ , and let  $\alpha$  be a variable of type  $e$ . Then  $\text{every}(\lambda \alpha / s \vdash K_1 / \lambda \alpha \wedge K_2)$  is a meaningful expression of type  $t$ , and  $\|\text{every}(\lambda \alpha / s \vdash K_1 / \lambda \alpha \wedge K_2)\|_{M, g, d, c, C}$  is the function  $f(s)$ :

$$f(s) = \begin{cases} \text{undefined} & \text{iff either } \|\lambda \alpha / s \vdash K_1\|_{M, g, d, c, C} \text{ is undefined or } \|\lambda \alpha / s \vdash K_2\|_{M, g, d, c, C} \text{ is undefined,} \\ 1 & \text{iff for every } g \text{ s.t. } \|\lambda \alpha / s \vdash K_1\|_{M, g, d, c, C}(g) = 1, \|\lambda \alpha / s \vdash K_2\|_{M, g, d, c, C}(g) = 1, \text{ where } K_1 \cup K_2 \text{ is the DRS whose markers are the union of the markers of } K_1 \text{ and } K_2, \text{ and whose conditions consist of all the conditions of } K_1, \text{ in their order, followed by all the conditions of } K_2, \text{ in their order} \\ 0 & \text{otherwise} \end{cases}$$


It is easy (although rather tedious) to check that this definition ensures that the discourse marker  $y$  occurring in the nuclear scope of (4.7) is defined, although it does not get bound by the occurrence of  $y$  in the restriction.<sup>20</sup> (i.e. it is possible that the value assigned to  $y$  to satisfy the restriction is not the same as the value assigned to  $y$  to satisfy the nuclear scope). The definition can be generalized to the case in which  $\|\text{every}(\lambda \alpha / s \vdash K_1 / \lambda \alpha \wedge K_2)\|_{M, g, d, c, C}$  is set-valued in the same way than the simpler definition above was.

I will note here that the definition above gives us what Chierchia calls the weak reading of the determiner 'every'.<sup>21</sup> Since it is only the farmers that are universally quantified over in order for (4.7) to be true at a situation, it is not necessary that a farmer beat all the donkeys he owns. This rescues the formalism from the so-called *proportion problem* illustrated by sentences such as "Most farmers who own a donkey thrive".<sup>22</sup> On the other hand, it is felt by some that the other interpretation should be available as well. I won't discuss this issue any further here.


Some might feel that the simpler representation in (4.8) might be adequate for the sentence "Every kid climbed a tree".

<sup>20</sup> This definition has nothing to do with the distinction between strong and weak NPs discussed in Chapter 3.

<sup>21</sup> These sentences are felt not to require quantification over farmer-donkey pairs in the same way as sentences such as "Every farmer who owns a donkey beats it." (Kadmon, 1997; Chierchia, 1999).



Indeed, I use this simpler representation as a shorthand in most cases of quantification discussed in the discussion. That the more complex (4.6) is needed is, however, shown by examples such as "Every school sent the principal to the meeting", that has one (preferred, for some readers) interpretation in which the resource situation for the definite "the principal" is the 'case' associated with the school. With the 'more complex' representation, this reading can be represented as shown in (4.9).



It can also be shown that the CRT translation for quantifiers assigns the proper truth conditions to sentences such as "No student who was told kind words by a teacher forgot her when he became an admiral.", in which the situation at which the nuclear scope is evaluated need not be the same as the situation in which the restriction is evaluated.

Finally, the 'Epistemic Logic-like' representation for quantification used in the rest of the dissertation for the lexical semantics of quantified expressions can be defined as follows:

$$\| \text{every}(\lambda x / s \vdash \Phi) \|_{\Psi} \equiv_{\text{def}} \text{every}(\lambda x / s \vdash \Phi) / \lambda x \wedge \Psi$$


And the same is true for the other determiners.

Weak NPs are ambiguous between a presuppositional reading, represented as above, and a non-presuppositional reading. The non-presuppositional reading of cardinal NPs can be represented by means of discourse markers and by assuming a treatment of plurality along the lines of Link.

#### 4.3.9 Conditionals and Adverbs of Quantification

Although adverbs of quantification are not included in the TRAINS fragment, I will briefly discuss their semantics in CRT for completeness' sake. Adverbs of quantification like "always" are treated in CRT as quantifiers over situations, as in Ide Swart, 1991; Stump, 1983; Berman, 1987; Chierchia, 1992; Heim, 1990. The translation of (4.10) in CRT is shown in (4.11)

(4.10) If a farmer owns a donkey, he (always) beats it



The non-set valued definition of  $\|K_1\| \xrightarrow{\text{always}} K_2 \|M.g.d.c.C\|$  is as follows

• Let  $K_1$  and  $K_2$  be DRSS. Then  $K_1 \xrightarrow{\text{always}} K_2 \|M.g.d.c.C\|$  is a meaningful expression of type  $t$  and

$\|K_1\| \xrightarrow{\text{always}} K_2 \|M.g.d.c.C\|$  is the function  $f$  such that

$$f(s) = \begin{cases} \text{undefined} & \text{if either } \|K_1\| \|M.g.d.c.C\| \text{ or } \|K_2\| \|M.g.d.c.C\| \text{ is undefined} \\ \text{over } s \\ 1 & \text{iff for each sub-situation } s' \text{ of } s \text{ such that } \|K_1\| \|M.g.d.c.C\| = 1 \text{ there} \\ & \text{is a substitution } \eta \text{ s.t. } s' \subseteq s'' \text{ s.t. } \|K_1 \cup K_2\| \|M.g.d.c.C\| = 1 \\ 0 & \text{Otherwise} \end{cases}$$

As the case  $C'(\|K_1\| \|M.g.d.c.C\|)$  extends  $C'(\|K_2\| \|M.g.d.c.C\|)$ , it is defined over the discourse markers  $x$  and  $y$ , and assigns to them the same values that get assigned to them by  $C'(\|K_1\| \|M.g.d.c.C\|)$

The definition can be expanded into a definition for the set-valued version of  $\|K_1\| \|M.g.d.c.C\|$  in the same way that the clause for every is


• Let  $K_1$  and  $K_2$  be DRSS. Then  $K_1 \xrightarrow{\text{always}} K_2 \|M.g.d.c.C\|$  is a meaningful expression of type  $t$ , and  $\|K_1\| \xrightarrow{\text{always}} K_2 \|M.g.d.c.C\|$  is the set of functions  $\{f \mid \text{for } h \in \|K_1\| \|M.g.d.c.C\|, \text{ and } h' \text{ in } \|K_2\| \|M.g.d.c.C\|, f \text{ is the function such that}$

$$f(s) = \begin{cases} \text{undefined} & \text{if either } h(s) \text{ or } h'(s) \text{ is undefined} \\ 1 & \text{iff for each sub-situation } s' \text{ of } s \text{ such that } h(s) = 1, \text{ there is a} \\ & \text{substitution } \eta' \text{ of } s' \subseteq s'' \text{ s.t. } h(s'') = 1 \\ 0 & \text{Otherwise} \end{cases}$$

### 4.3.1b Modals

Modals such as 'must' and 'should' are also treated in CRT as quantifiers over situations. I repeat here the CRT representation for the sentence 'An engine must go to Avon' discussed in Chapter 1:

(4.12) An engine must go to Avon



The semantics of  $K_1 \xrightarrow{\text{must}} K_2$  is inspired from the approach taken in Disc Semantics II, and in (1986a), where the semantics of modal statements is formulated in terms of situation extensions and makes without accessibility relations. The result is that the semantics of must is similar to the semantics of always

• Let  $K_1$  and  $K_2$  be DRSS. Then  $K_1 \xrightarrow{\text{must}} K_2$  is a meaningful expression of type  $t$ , and  $\|K_1\| \xrightarrow{\text{must}} K_2 \|M.g.d.c.C\|$  is the set of functions  $\{f \mid \text{for } h \in \|K_1\| \|M.g.d.c.C\| \text{ and } h' \text{ in } \|K_2\| \|M.g.d.c.C\|, f \text{ is the function such that}$

$$f(s) = \begin{cases} \text{undefined} & \text{if either } h(s) \text{ or } h'(s) \text{ is undefined} \\ 1 & \text{iff for each sub-situation } s' \text{ of } s \text{ such that } h(s) = 1, \text{ there is a} \\ & \text{substitution } \eta' \text{ of } s' \subseteq s'' \text{ s.t. } h'(s'') = 1 \\ 0 & \text{Otherwise} \end{cases}$$

This semantics of modals is motivated in greater detail in §5.9

#### 4.4 DRS CONSTRUCTION RULES AS INFERENCE RULES

In this section I discuss in more detail the behavior of the construction rules used to formalize discourse interpretation in CRT. Although these construction rules are very much like those used in DRT, their application may result in conflicting hypotheses, this aspect of the theory is clarified a bit in this section. I also attempt to show the connection between this form of inference and the traditional notion of inference formulated in terms of inference rules, and in particular with the formulation of non-monotonic reasoning due to Reiter (Reiter, 1980). This part of the proposal is, however, very preliminary.

It would be of interest to study the formal system thus defined in some detail, and providing a semantics for logical forms is a first step towards a system in which some construction rules can be semantically justified. However, I concentrate in this dissertation on construction rules that are defeasible and therefore all I am going to do in this section is to provide a sketchy description of how the effect of construction rules, and especially defeasible construction rules, can be formalized. I leave a more detailed study of the formal system for further research.

##### 4.4.1 Construction Rules as Operations on DRSs

Inference is traditionally formalized in terms of (additive) operations on sets of formulas called inference rules. The construction rules that get applied in the DRS construction algorithm can be thought of as a generalization of the inference rules found in traditional logic. Whereas an inference rule in first order logic (say, Modus Ponens) is an operation on sets of formulas a DRS construction rule is an operation on two sets: a set of discourse markers and a set of conditions (that are essentially formulas). That is, where's an inference rule in traditional logic is an operation that takes a set of formulas as input and returns a set of formulas as output, i.e., of the form

$$\text{IR: } \varphi(L) \rightarrow \varphi(L)$$

(where  $\varphi(L)$  is the powerset of formulas of the language  $L$ ), a DRS construction rule is an operation that takes a DRS as input and returns a DRS as output, of the form


$$\text{CR: } \langle \varphi(M) \times \varphi(L) \rightarrow \varphi(M) \times \varphi(L) \rangle$$

Where  $\varphi(M)$  is the powerset of the set of marks:  $\mathbf{i}, \mathbf{M}$ .

A second sense in which the construction rules used in DRT and CRT generalize the inference rules of traditional logic is that inference rules are always *additive*, in the sense that the 'output set' of an inference rule always contains the 'input set' of formulas, whereas a construction rule in DRT typically operates by *replacing* a formula in the input DRS, that is, by *deleting* some formula and *adding* a new formula that is less underspecified than the input formula. In this sense, the construction rules resemble the operations allowed in the belief revision literature (Gärdenfors, 1988; Nebel, 1990; Gärdenfors, 1992).

The DRS construction rules are combinations of additions and replacements on the marker set and condition set that make up a DRS. Consider, for example, the DRS construction rule

EXAMPLE CR below, a fairly typical example of the kinds of DRS construction rules discussed in Chapter 6.



This rule can be thought of as an operation that can apply to a DRS  $K$  if  $K$  contains among its conditions an underspecified representation matching the trigger and the two conditions  $\Phi$  and  $\Psi$ . The condition set resulting from the inference rule is obtained by deleting the trigger from the initial condition set, and adding one marker  $z$  and two conditions: the logical form that replaces the trigger, and the condition  $\Xi$ .

An inference rule like Modus Ponens can clearly also be defined as an operation on DRS, although one that does not add any new markers and does not replace its trigger.

MODUS PONENS	
Triggering configuration	A
Constraints	$A \rightarrow B$
Replace with:	B
Introduce into Context:	

Having established a connection between the notion of construction rule and the notion of inference rule, we can obtain a notion of 'derivability'—this is a relation between DRSs defined with respect to a set  $\Delta$  of construction rules

**Definition 4.8.** The DRS  $K'$  follows from the DRS  $K$  with respect to the construction rules set  $\Delta$ ,  $K' \vdash_{\Delta} K$  if  $K'$  is consistent, it is either the same as  $K$ , or it is obtained by applying a construction rule in  $\Delta$  to a DRS  $K''$  such that  $K \vdash_{\Delta} K''$

I will call a DRS  $K$  such that  $K' \vdash_{\Delta} K$  an hypothesis derived from  $K'$  by  $\Delta$

It's also useful to define the notion of complete hypothesis of a DRS  $K$  under the set of construction rules  $\Delta$ , written  $CH_{\Delta}(K)$  a complete hypothesis is an hypothesis  $K'$  derived from  $K$  by  $\Delta$  that is 'maximal' in the sense that no construction rule in  $\Delta$ , when applied to  $K'$ , results in a larger consistent hypothesis derived from  $K$  by  $\Delta$ ,  $K''$ .

Just as in first order logic, a complete hypothesis derived from  $K$  by  $\Delta$  in general will not be finite—it is enough for a complete hypothesis to be infinite that the 'initial hypothesis'  $CON_K$  contain at least two conditions, and that  $\Delta$  contains a rule of conjunction introduction. It's also clear that a DRS  $K$  will result in a single complete hypothesis under  $\Delta$  only if all construction rules in  $\Delta$  are sound, more in general, there will be more than one complete hypothesis.


The notion of complete hypothesis is in part derived from Reiter's notion of extension (Reiter, 1980). It is known that, in general, a default theory doesn't always have extensions (extensions being the closest notion in Default Logic to the notion of closure). It's not clear yet whether the argument for the lack of extensions in DL transports to CRT, but in the system sketched here there is an additional reason for DRSs not to originate complete hypotheses under a set of construction rules  $\Delta$ . This is because a construction rule may replace a condition in a DRS, instead of simply adding new conditions to it. Thus, for example, if  $K$  is a DRS that

contains only one condition,  $A$ , and if  $\Delta$  consists of two construction rules—one that replaces  $A$  with  $B$ , the other that replaces  $B$  with  $A$  again—there will be no complete hypothesis derived from  $K$  under  $\Delta$

I will add, however, that this is not a problem for the class of construction rules we consider. All the construction rules used in this dissertation that actually include a replacement operation are such that they replace an underspecified representation with a less underspecified one. In other words, they are all of the form 'Replace  $A$  with  $A'$ , where  $A' \vdash A$ . It's easy to see that this cannot lead to an infinite 'flip-flop'. (Note that things would be different were we to allow for arbitrary deletions.)

It is important to understand that because I am only interested in characterizing the set of inferences that actually get made during surface discourse interpretation, the set  $\Delta$  of construction rules does not include the set of inference rules of first order calculus, thus the notion of 'deductive closure' and of 'extension', that are distinct in Reiter's system, are merged here

In order to simplify the formulation of the construction rules, I allow some syntactic sugar in the specification of the trigger. Where Kamp and Reyle would have disjunctions in the trigger of their conditions—that is, they would have construction rules that allow more than one trigger (and can be thought of as corresponding to a set of construction rules) I will use triggers of the form



#### 4.4.2 Weak and Strong Inference Rules

In the system just sketched, all construction rules have the same 'strength'. As discussed in §1.5, however, it appears that in discourse interpretation some rules may override others. I proposed in §1.5 a distinction between two kinds of construction rules: *strong* ones and *weak* ones. I also proposed that the hypotheses generated by discourse interpretation are then ranked on the basis of their plausibility. Accordingly, we want to characterize the set of hypotheses obtained by discourse interpretation in terms of a discourse inference system: a triple:

$$D = (K, \{\Delta_1, \Delta_2, \dots, \Delta_n\}, \pi)$$

where  $K$  is a root DRS,  $\Delta_1, \Delta_2, \dots, \Delta_n$  are  $n$  sets of construction rules, and  $\pi$  is a plausibility function. The set of complete hypotheses with respect to a discourse inference system  $D$  is defined as follows. The construction rules are arranged in sets in such a way that the inference rules in  $\Delta_j$  take precedence over the inference rules in  $\Delta_i$  if  $j < i$ . 'Taking precedence' here is meant in the sense in which a set of beliefs takes precedence over another in Hierarchic Autoepistemic Logic (Appelt and Konig, 1988). Each set of construction rules applies only to the consistent complete hypotheses resulting from the application of all the higher-priority sets of construction rules. First the construction rules at level 1 are applied to initial DRS  $K$ . Because conflicts

may result from the application of these constructive rules, in general, more than one consistent complete hypothesis may be found. Then, the inference rules at level 2 are applied to each of these complete hypotheses, and so forth.<sup>27</sup>

More formally, the set of complete hypotheses under a discourse inference system  $DHS$  is defined as follows. Let  $H_{\infty}$  be the DHS  $K$ , and let  $H(S, n) = \{H_{\infty}\}$ . Then, for  $j = 1, \dots, n$ ,  $H^j(S)$  is the set of all consistent complete hypotheses  $\{H_1^j, \dots, H_m^j\}$  derivable from a complete hypothesis  $H_{j-1}^j$  in  $H^j(S, j-1)$  under a subset of the axioms in  $(\Delta_1, \Delta_2)$ . The resulting set of complete hypotheses  $H^j(S) = H^j(S)$ .

Once all set of extensions of the discourse inference system has been computed, the plausibility function  $\lambda$  is invoked to determine a partial order between these extensions. A preference ambiguity is observed if  $H^j(S)$  has more than one element, and there is no complete hypothesis  $H_{j-1}^j$  such that  $H_{\infty}$  is 'more plausible' than all of the other elements in  $H^j(S)$ .

In practice, only two levels of default inference rules are used in Chapter 6—I'll call them strong defaults and weak defaults, respectively.<sup>28</sup>

#### 4.4.3 Syntactic and Semantic Definitions of Inference

Although the construction rules are defined in a syntactic fashion—as operations on sets of formulas—a connection between the syntactic definition and the semantics of the language is not precluded. That is, it is possible to characterize these operations as operations on sets of (sets of) propositions.

For example, it is possible to characterize an operation as sound if it does not result in an interpretation which was not included in the space of possibilities left by the original hypothesis, thanks to the fact that CBT underspecified representations are assigned a semantic interpretation. I also discussed in §3.8.6 how I propose to account for the effect on interpretation of syntactic factors such as PP preposing by assuming that the common ground includes information about the syntactic structure of utterances, in the form of 'micro conversational events'.

<sup>27</sup> Note that in general there is an recursive procedure that constructs the sets of complete hypotheses at level 1, etc., each of these sets may in general be infinite.

<sup>28</sup> There seems to be an interesting connection between this way of defining preferences among hypotheses and Brevik's proposal for improving a preference system over Brevik's framework for default reasoning (Brevik, 1988) in chapter 5 of [Brevik, 1988]. The notions of 'strong plausibility' and 'weak plausibility' defined there, however, are mostly unrelated to the notions of 'strong default' and 'weak default' used here.

## 5 Syntactic Structure and Lexical Semantics in a TRAINS fragment

The scope assigned to operators is in part determined by syntactic and semantic constraints, in part by the result of discourse interpretation procedures. In this chapter I present my semantic analysis of several classes of lexical items, together with my analysis of a number of constructions with interesting properties as far as scope is concerned, and I discuss how this syntactic and semantic information results in scopable preferences such as the Quantifier Hierarchy. The 'fragment' or subset of English, defining the cases of scope assignment that I analyze is given in large part determined by the TRAINS corpus. This 'TRANS fragment' includes a subset of the operators listed in §1.4 that includes definite and indefinite descriptions, modals, tense and the quantifier 'every'. The TRAINS fragment also includes a variety of linguistic constructions such as temporal adverbials and there-insertion clauses.

Each section of this chapter is concerned either with a class of lexical items (e.g., proper names, definite descriptions or modals) or with a particular construct, such as there-insertion sentences or infinitival clauses. Both the semantic translations of lexical items and my assumptions about syntax are discussed; the semantics of the logical forms fragments involved with each construction is also provided.

### 5.1 NOUN PHRASES THAT ARE NOT OPERATORS: PROPER NAMES, PRONOUNS, AND BARE PLURALS

I will first of all consider those noun phrases that are not operators in the sense discussed in §3.2.1.e, that do not originate scopable ambiguities. As discussed in §3.2, these are the NPs whose logical form has a Proper Value that can be specified without using the strategic mechanism. There are three such cases of NPs in the TRAINS fragment: proper names, pronouns, and bare plurals.<sup>1</sup>

#### 5.1.1 Proper Names

Proper names are translated as constants of type  $e$ . For example, in the proper name 'Cernig' has the following translation:

<sup>1</sup> There are other such NPs in the TRAINS conversations, most notably demonstratives, but I will not discuss these.

"Coming"  $\sim x$

Here,  $x$  is a constant of type  $e$ . I assume that type-raising operations [Parice and Rovak, 1987] are available to "lift" the translation of proper names to that of quantified NPs, so as to obtain a denotation that can be compared with that of a quantified NP, such as "several of his friends" as in 'John and several of his friends'.

### 5.1.2 Pronouns

First and second person pronouns are translated using the indexicals *self* and *other* discussed in §4.3. A small complication is the fact that "I" refers to *self* during generation, and to *other* when interpreting an utterance of the other speaker  $s$ , so in fact "I" should be translated as a parameter that may assume one of these two values. I ignore this complexity here, and simply assume that special indexicals *I* and *we* exist that are used to translate these pronouns and get interpreted appropriately.

- "I"  $\sim I$
- "You"  $\sim you$
- "We"  $\sim we$

Third person pronouns are referentially ambiguous expressions, and are therefore translated as parameters of type  $e$ , as already discussed in §3.4 and §4.3. This also makes pronominal presuppositional NPs, which seem correct (cf. Heim's discussion in [Heim 1987]).

- "He"  $\sim x$

This translation, while adequate for my purposes, falls short in two respects. First of all, proper names should be classified as presuppositional NPs. In §3.6 I proposed to reformulate the classification of NPs as presuppositional and non-presuppositional to make it depend on whether the NP's classification was *participative* or *non-participative*, but the translation above makes proper names non-presuppositional. Secondly, proper names in our conversation may or may not have a predicative component, as in "Engine 1" or "city of Rome". This requires a translation closer to those used in DRT, i.e. of the form  $\lambda x.P(x) \wedge P(x)$  where  $x$  is a discourse marker, except with the discourse marker replaced by a parameter as follows:

$$x \text{ "Coming"} \sim \lambda x.P(x) \wedge P(x)$$

Kamp et al. (1990) show a way to "unburden" translations of this kind so that one would get a referential interpretation along the lines of Kratke's theory of reference [Kratke 1972].

Again, a more complex translation would arguably be required taking into account gender and number of the pronoun, such as:

$$x \text{ "he"} \sim \lambda x.P \text{ MALE}(x) \wedge \text{ANY}(x) \wedge P(x)$$

However, it is not at all clear whether this kind of information should be part of the semantics instead of being represented in the form of syntactic features. This would also be part of the logical form and thus could still be used to trigger the appropriate interpretation rule. At the moment, I assume the latter solution.

- "him"  $\sim x$
- "her"  $\sim x$
- "it"  $\sim x$
- "they"  $\sim x$

This representation can be used for referential as well as anaphoric, and  $e$  type uses of pronouns.

### 5.1.3 Bare Plurals

Fredriksson Carlson, I assume that bare plurals such as "oranges" are uniformly translated as kind-denoting terms. As discussed above, I also assume that kinds are predicate nominalizations, as suggested by Ciccharelli, Cherchia, and Schubert and Pelletier. The translation for the bare plural "oranges" is shown below.  $K$  is the kind forming operator proposed by Hwang and Schubert.

$$\text{"oranges"} \sim K(\text{ORANGE})$$

### 5.1.4 Logical Forms for Non-operator Noun Phrases

The syntactic analysis of noun phrases is the subject of much debate these days (e.g., [Abney, 1987]). I omit all discussion and simply assume, in this section as well as in the sections dedicated to the semantic translation of quantified NPs, that NPs are primitive constituents of logical form, in that the logical form fragment associated with NP is of the form:

$$(S) \quad [np \ NP]$$

where NP is the translation obtained by combining the translations of the subcomponents, if any. E.g., the logical form fragment for the proper name *Coming* whose semantic translation was shown before is of the form

$$[np \ c]$$

and one logical form for the pronoun "he" is as follows

$$[np \ x]$$

The denotation of these logical form fragments depends on the kind of noun phrase. As none of the NPs discussed in this section is an operator, the Cooper Value CV of the logical forms associated with these NPs is always a singleton set. If  $[np \ NP]$  is such a logical form fragment, then

$$CV([np \ NP]) = \{ (NP) \}$$



This means that there are no alternative ways of 'using' the meaning of one of these NPs in a derivation, and therefore no scope ambiguities may result from the interaction of these NPs with other NPs, although the resulting interpretation may still be ambiguous —e.g., in the case of pronouns that denote parameters.<sup>4</sup>

## 5.2 EVERY-NPS AND OTHER QUANTIFIERS

### 5.2.1 The Semantic Translation of Strong and Weak Determiners

As discussed in §3.6 and §4.3, I adopt Milnik's distinction between strong and weak NPs, formulated as a distinction between NPs with existential presuppositions and NPs without existential presuppositions, and I assume that an existential presupposition is a statement of the form  $\exists x K$ , where  $x$  is a parameter of type  $s$ .

NPs of the form "det Q" where "det" is a strong determiner, such as "every/lexical", are assigned the schematic translation in (5.2)

$$(5.2) \quad \text{"det Q"} \sim \lambda P [\text{DET } x / s \models Q(x)] (P(x))$$

I discussed in §4.3 that (5.2) is shorthand for the more complex, generalised quantifiers' translation

$$\text{every}(\lambda x / s \models \Phi) (\lambda x \Phi)$$

Weak determiners, such as "many" or "some", are ambiguous between a strong interpretation (represented as in (5.2)) and a weak reading, represented as follows

$$(5.3) \quad \text{"det Q"} \sim \lambda P [\text{DET } x / Q(x) \wedge P(x)]$$

The semantics of these translations is discussed in §4.3.

### 5.2.2 Logical Forms

As mentioned in the previous section, I assume that NPs are "basic" —i.e., undecomposed— logical form constituents. This also applies to quantified NPs, which means I will not discuss here the syntactic structure of relative clauses and other prenominal or postnominal modifiers (There are very few cases of complex NPs in the TRAINS conversations). The logical form translation associated with both strong and weak NPs is of the form

$$[\text{NP NP}]$$

<sup>4</sup>This claim seems to be rather uncontroversial in the case of pronouns and proper names, even in the case of bare plurals. It has been argued that bare plurals are ambiguous between an existential and a kind reading (Dyck, 1992). I ignore this possible complexity here.

where NP' is the semantic translation of the NP. I treat presuppositional quantifiers as operators— that is, I assume that the logical form fragment associated with a presuppositional quantifier has a Cooper Value consisting of two distinct sequences, one of which is the semantic translation of the quantifier, the other obtained by putting the quantifier 'in storage'. In order to make it simpler to define the 'storage discharging' operations later, I propose that what gets put in storage is not the translation of the quantified NP, but the function from propositions to propositions obtained by applying the translation of the quantified NP to the predicate  $\lambda x \Phi$ , where  $\Phi$  is a variable of type  $t$ , and then abstracting over  $\Phi$ ; this results in a uniform storage value for tense, modals, and quantified NPs

$$\bullet \text{ CV}([\text{NP } \lambda P [\text{DET } x / s \models Q(x)] (P(x))]) = \left\{ \left( \lambda P [\text{DET } x / s \models Q(x)] (P(x)) \right), \left( \lambda x \lambda \Phi [\text{DET } x / s \models Q(x)] (\Phi) \right) \right\}$$

I assume instead that the non-presuppositional interpretation of weak NPs does not result in an operator; that is, that the logical form associated with that reading has a 'Cooper Value' that consists of a single sequence:

$$\bullet \text{ CV}([\text{NP } \lambda P [\text{DET } x / Q(x) \wedge P(x)]]) = \{ (\lambda P [\text{DET } x / Q(x) \wedge P(x)]) \}$$

This distinction is related to the hypothesis, most recently advanced by Diesing (1992), that the non-presuppositional interpretation of weak NPs is not subject to Quantifier Raising (in my terms, it is not an operator) and therefore only weak NPs get wide scope on the presuppositional and the specific reading.<sup>5</sup> What in Diesing's proposal is a syntactic property of these NPs (i.e., whether or not they undergo Quantifier Raising) becomes now the consequence of three hypotheses: that the aim of discourse interpretation is to assign values to the parameters that occur in a logical form; that scope assignment is driven by this process of discourse interpretation, in the sense that DRT-style 'model construction rules' are activated when these parameters are resolved; and that there are two classes of NP interpretations, some that require parameters to be identified, and others that do not contain parameters. If we put these hypotheses together, we get the result that unless the semantic translation of an NP is parametric, no model construction rules are applied, and therefore the NP is interpreted in the scope of whatever operator was part of the sentence; in effect, the NP 'doesn't raise'.

We also get the prediction that the easier it is for an NP to be interpreted as presuppositional, the stronger its tendency to take wide scope. This is because there are no discourse interpretation rules triggered by the non-presuppositional reading of weak NPs, since this reading does not contain any parameters. This might explain some of the Quantifier Hierarchy effects observed in the literature. Van Lehn already observed something of the sort and proposed a 'definiteness hierarchy'. More in general, the tendency of indefinite NPs and other weak NPs in object position to take narrow scope may result from a difficulty of interpreting them as presuppositional in that context.<sup>6</sup> I agree with van Lehn that the 'Quantifier Hierarchy' is in fact the result of two

<sup>5</sup>For Ex, there are two readings available

<sup>6</sup>Although the choice of an interpretation for a weak NP is, strictly speaking, a matter for a discourse or lexical disambiguation, it seems a bit odd to have on the same entity. I believe that Diesing's hypothesis that the syntactic position affects the choice of interpretation is on the right track. Namely, a weak NP that is in a syntactic position marked as part of the 'new' information (inside a VP) tends to be interpreted as non-presuppositional, whereas a weak NP that is outside of VP tends to be interpreted as presuppositional.

distinct factors: context-dependency and distributivity, and I discuss the scoping differences between 'singular' universal NPs such as "every boxer" and plural definite NPs with universal force such as "the boxers" below.

I must add that the fact that non-presuppositional interpretations of weak NPs do not behave as operators has been established beyond any doubt, so it should be taken as just a reasonable hypothesis. Two issues that need to be addressed are whether the presuppositional and specific reading of indefinites are, in fact, the same, and whether NPs like "a boxer" count as weak, which they should according to Millard's there-insertion test. Assuming that the answer to both questions is yes, we need to ask whether "a kid" in (5.4a) and "a banana peel" in (5.4b), both of which seem to take wide scope on the preferred reading of the sentence in which they occur, are presuppositional. I would say that perhaps a case may be made for (5.4a), but I'm more skeptical about (5.4b). I'll leave this issue for further research.

- (5.4) a A kid climbed every tree.  
b John didn't see a banana peel and slipped.

## 5.3 DEFINITE DESCRIPTIONS

### 5.3.1 The Semantics of Definite Descriptions: Two Proposals

Two main proposals exist concerning the semantics of definite descriptions. Russell (1905) proposed that a sentence like "The dog came in" asserts that there is a dog, that this dog is unique, and it came in. This translation can be represented as in (5.5).

- (5.5) The dog came in.  
(5.6)  $\exists x \text{ DOG}(x) \wedge (\forall y (y = x)) \wedge \text{CAME-IN}(x)$

Three aspects of this proposal have been criticized. Strawson (1950) argued that sentence (5.5) is neither true nor false when there is no dog, the existence of a dog, according to him, is not asserted by (5.5), but only presupposed. Several authors have argued that definite descriptions are not quantificational, perhaps the best argued of these theories is Heim's (1982), that I discuss below. Finally, it has been observed (e.g. by Lewis (1979)) that uniqueness needs to be qualified somehow, or else a sentence like (5.5) could only be true in a world that includes a single dog.

The Russellian proposal is nevertheless enjoying renewed interest these days, thanks to the work of Grice (1969), Kaplan (1977), Neale (1990) and Kadmon (1987). According to Grice, it is necessary to distinguish between what the speaker says and what the speaker means. The truth conditions of an utterance of a sentence of the form 'the F is G' are thus strictly Russellian: the hearer that a particular individual *b* is G, and may succeed in doing this by exploiting the fact that both speaker and hearer take *b* to be the F (Neale, 1990, p. 9). So, while "it is surely not open to dispute that a sentence of the form 'the F is G' may be used to communicate an object-dependent thought to someone, to the effect that some particular individual *b* is G" (Neale, 1990, p. 7), "in our daily talk we very often convey things" indirectly, relying on what we take to be our interlocutors' abilities "to grasp what we mean by our utterances" (Neale, 1990, p. 9).

Neale implements this proposal by introducing in the object language a determiner, THE, defined as in (5.7), and by assuming that definite NPs such as "the dog" have a unique translation of the kind in (5.8) (Neale, 1990, p. 45):

- (5.7)  $\forall \text{THE } x \text{ F}(x) / \text{C}(x) \text{ is true iff } [\text{F} - \text{G}] = 0 \text{ and } |\text{F}| = 1$   
(5.8) "the dog"  $\sim \lambda Q \text{ THE } x \text{ DOG}(x) / Q(x)$

Lewis (1979) and Kadmon (1987) also suggest that uniqueness can be rescued if we assume that domain selection takes place, i.e., that definite descriptions are just like all other quantified NPs in that they meet the set of quantificational to be contextually determined.

An alternative approach to the semantics of definite descriptions was proposed by Heim in her dissertation (1982). According to Heim, definite and indefinite NPs are not quantifiers, but open variables. The truth conditions for (5.5) proposed by Heim are as follows:

- (5.9)  $\text{DOG}(x) \wedge \text{CAME-IN}(x)$

in (5.9), *x*, is an open variable. Furthermore, Heim proposes that the sentence "a dog came in" has the same translation, the difference between definites and indefinites is that (i) definites must be familiar in the context and (ii) the restriction in a definite NP in the case of "the dog" the predicate  $\text{DOG}(x)$  must be presupposed.

These aspects of the interpretation of definites and indefinites are captured with reference to the common ground. Heim argues that the right metaphor for thinking about the common ground is that of *file cards* which are augmented as a discourse proceeds with the content of each utterance, and whose domain consists of the indices of definite and indefinite NPs. Heim then characterizes the difference between definites and indefinites in terms of a Novelty/Familiarity Condition that reads as follows (Heim, 1982, p. 370).

#### Novelty-Familiarity-Condition:

- For  $\Phi$  to be felicitous w.r.t. a file F it is required for every NP, in  $\Phi$  that  
i if NP, is {definite}, then i  $\notin \text{Dom}(F)$ ;  
ii if NP, is {indefinite}, then  
a  $i \in \text{Dom}(F)$ , and  
b if NP, is a formula, F entails NP.

Heim argues that the fact that definite descriptions refer uniquely is not a consequence of their semantic interpretation, but of the fact that the listener, upon hearing a definite NP, has to identify a unique referent for it in the common ground, which cannot be done unless the common ground contains only one appropriate object.

There are good arguments both in favor of the Russellian position and in favor of Heim's theory, and the debate about the relative merits of these two positions has been extensive. The proposals about definite description interpretation I present in §6.5 do not really depend on choosing one or the other framework; however, I believe that Heim's approach fares better than the Russellian proposal in the case of definite NPs of the form "the N of NP" like those in (5.10).

- (5.10) a I got these data from the student of a famous linguist (Zamparelli, p. c)

- b The village is located on the side of a mountain
- c I usually had breakfast at the corner of a major intersection
- d On Friday, a bomb exploded outside the offices of an American corporation

All of the definite NPs in (5.10) have an interpretation that does not require uniqueness in the sense of Russell (5.10a), for example, has an interpretation that can be paraphrased as "There is a famous linguist, and there is a student of this famous linguist, and this student gave me these data." This interpretation does not require either there to be a unique student of a famous linguist, or the famous linguist in question to have a unique student. (Similar examples were discussed by Lehrer (1987).) Although these examples are problematic for Heim's theory as well, in that no familiarity with the student in question seems to be requested, it is possible to modify the conditions under which familiarity is requested in such a way as to obviate the problem, as I discuss in [Poesio, 1994]. It is not possible, however, to reconcile these examples with a theory like Russell's, since in that theory uniqueness is the only difference between NPs like "the student of a famous linguist" and NPs like "a student of a famous linguist".

Before discussing my semantic translation for definite descriptions, I will discuss a theory of definite description interpretation that is closely related to Heim's theory, and that gives some insights into the appropriate translation for definite descriptions.

### 5.3.2 Interpretation of Definite Descriptions: The Location Theory

The *location theory of definite descriptions* (Hawkins, 1978; Clark and Marshall, 1981) is perhaps the most widely accepted account of the processes resulting in an hearer's assigning a referential interpretation to a sentence containing a definite NP. According to Hawkins (1978, p. 167) the defining aspects of this process are that

- 1 "... the hearer is instructed to locate the referent in some *shared set of objects*" (emphasis added)
- 2 "the speaker refers to the *totality* of the objects/mass within this set that satisfy this restriction" (emphasis added)

This theory can be seen as combining an hypothesis about the semantics of definite descriptions (point 2), with one about the process of pragmatic interpretation (point 1). According to the location theory, the aim of definite description interpretation is (i) to identify a shared set of objects, and (ii) to identify the referent of the definite within that shared set.

Clark and Marshall (1981) make the notion of "shared set" used by Hawkins more precise. They argue that what is needed is the notion of *mutual knowledge* introduced by Lewis and Harman (Lewis, 1969; Harman, 1977) among others. Clark and Marshall adopt the "recursive" definition of mutual knowledge proposed by Harman (1977), cited by [Clark and Marshall, 1981] p. 17)

- (5.11) A and B mutually know that  $p \equiv \text{def}$  (i) A and B know that  $p$  and that  $q$

Clark and Marshall also explain how people can infer mutual knowledge, which apparently requires checking an infinite amount of conditions of the form "A knows that B knows that

that  $p$ ". Their solution is based on a proposal by Lewis: if A and B make certain assumptions about each other's rationality, they can use certain states of affairs (*grounds*) as a basis for inferring the infinity of conditions all at once. In the case of a "visible situation" use of definite descriptions, for example (see §3.7.1), the grounds consist of two parts: direct visual evidence of co-reference, and assumptions about the situation—that the other participant in the conversation is consciously attending, that he is rational, and so forth.

Heim's theory of definites and the location theory, under the theory of presupposition presented in §3.6, turn out to be closely related. I proposed in §3.6 that  $\Phi$  existentially presupposes  $\Psi$  if  $\Phi$  is of the form "... $x$ : $\Psi$ ..." that is, if in order to interpret  $\Phi$  it is required that the common ground contains a situation  $s$  that supports  $\Psi$ . It's clear that if we take the common ground to represent what's mutually known, then the situation that is needed to satisfy the presupposition is a "shared situation" in the sense of Hawkins. And to require that the referent of the description can be 'identified' in the shared set is tantamount to requiring that its index be identified as that of an existing file card. I discuss in the next section a semantic translation for definite descriptions that draws on both theories, while at the same time providing an indication of which aspects of the interpretation of a definite description have to be resolved by discourse interpretation processes.<sup>7</sup>

### 5.3.3 A Situation-Theoretic Formulation of the Semantics of Definite Descriptions

Taking Heim's theory as a starting point—in the sense that definite NPs are not translated as quantifiers, but as introducing discourse markers—and formulating the notions of 'presupposition' and 'familiarity' as suggested by the location theory, we get the translation for definites in (5.12). According to (5.12), "the boxcar" denotes the set of properties  $P$  that hold of a discourse marker  $x$  which has the property of being equal to an object  $y$  that is a boxcar in situation  $s$ . The interpretation of a definite NP contains two parameters: a parameter indicating that the resource situation of the NP has to be identified, much as in the case of strong quantified NPs, and a parameter indicating that the discourse marker has to be related to a previously introduced discourse marker. Following (Barwise and Perry, 1983, p. 145), I call the situation  $s$  which includes the objects quantified over by a determiner the *resource situation* in virtue of that determiner,  $s$  is the resource situation of the definite description in (5.12).

- (5.12) "the boxcar"  $\sim \lambda x P(x) \equiv [\text{BOXCAR}(x) \wedge x = y] / \wedge P(x)$

The location theory can be reformulated in terms of parameter anchoring as follows: the aim of the interpretation processes involved in definite description interpretation is to find an appropriate anchor for the situational parameter denoting the resource situation of the definite—i.e., to recognize an intention of the speaker's of the form:

$$\text{INTEND}(\text{anchor } s = s)$$

<sup>7</sup> Additional interpretive processes are known to be involved in the interpretation of definite descriptions (Clark, 1984), but in order to interpret the definite descriptions found in the current set of TRAINS transcripts all that the hearer is required to do is to identify a suitable "shared set" and its referent. Therefore, I am only concerned with this process in this discussion.

where  $S$  is the resource situation of the definite and  $s$  is the situation the speaker intends the hearer to 'locate' the referent of the definite in. The task of a theory of definite description interpretation, then, is to provide principles for anchoring resource situations that generate hypotheses about the intended identity of the resource situation of the determiner.

### 5.3.4 Logical Forms

The translation proposed above makes definite NPs string NPs, thus one would expect the associated logical forms to denote two distinct sequences, like other presuppositional NPs:

$$\bullet \text{ CV}(\text{INP } \lambda P / s \models \text{IQ}(x) \wedge x = y // \wedge P(x)) = \left\{ \begin{array}{l} (\lambda P / s \models \text{IQ}(x) \wedge x = y // \wedge P(x)), \\ (s \lambda \Phi / s \models \text{IQ}(x) \wedge x = y // \wedge \Phi) \end{array} \right\}$$

because, however, the value assigned in the discourse marker  $x$  only depends on the value of the parameter  $y$ , the interpretation of the definite NP is only going to be affected by other operators if the definite is interpreted as anaphoric on markers introduced in the scope of these operators, that is if the second parameter in the interpretation of the definite is 'anchored' to a marker in the scope of an operator, as in the examples in (5.13).

- (5.13) a If a dog sees a cat, the cat meows  
b Each school sent the principal to the meeting

I propose therefore that definites are not operators i.e., the Cooper value of the logical form associated with a definite is as follows

$$\bullet \text{ CV}(\text{INP } \lambda P / s \models \text{IQ}(x) \wedge x = y // \wedge P(x)) = (\lambda P / s \models \text{IQ}(x) \wedge x = y // \wedge P(x))$$

### 5.3.5 Plural Definite NPs

I assume that, once the semantic distinction between singular and plural predicates is factored out, plural definite NPs are translated much as singular definite NPs.

As mentioned in §4.1, my treatment of plurals is borrowed from Link's [1983], according to whom the universe  $I$  is a semi lattice whose elements model groups of atomic individuals, and whose partial order relation models group inclusion. I'll also recall from that section that the predicate modifier \* maps a predicate  $P$  onto a predicate  $P^*$  whose extension are groups whose elements are in the denotation of  $P$ . For example, the extension of the predicate BOXCAR\* consists of all the possible groups of boxcars.<sup>8</sup>

This results in the following translation for the plural definite NP 'the boxcars':

$$(5.14) \text{ "The boxcars" } \sim \lambda P / s \models \text{BOXCAR}^*(x) \wedge x = y // \wedge P(x)$$


<sup>8</sup> Link's treatment has been criticized by Landman [1989a, 1989b], who argues that there is no need for groups and 'multiples' can be used instead. Adopting Landman's multiple instead of Link's would not seriously affect my proposal.

What I said above about the Cooper Value of logical forms like [NP def], where 'def' is the translation of a singular definite NP, holds for plural definite NPs as well: the value is a singleton set consisting of the sequence {(def)}.

Notice that van Lehn's revision of Ioup's hypothesis about a quantifier hierarchy may be explained as a consequence of the difference between the translation of NPs such as "each child" and the translation of NPs such as "the children". Consider first the following revised version of Ioup's minimal pair ((1.44) discussed in Chapter 1)

- (5.15) a I saw a picture of each child.  
b I saw a picture of the children

((1.44a) has at least two semantically licensed translations, one in which the NP "each child" takes wide scope, the other in which the NP "a picture" takes wide scope, of the two, the first reading is preferred because "each child" only has a presuppositional interpretation, whereas, as I mentioned above and I discuss more in detail in the section below on indefinites, indefinite NPs tend to be interpreted non-presuppositionally unless they are in subject position or otherwise marked as specific (5.15b), however, has only one interpretation, that can be represented as in (5.16))



The contrast actually observed by Ioup (I repeat her minimal pair below), can be explained in the same way

- (1.44) a I saw a picture of each child.  
b I saw a picture of all the children

That "the children" and "all the children" are not completely synonymous is shown by the following contrast

- (5.18) a The children are numerous  
b \*All the children are numerous.

However, this contrast is not usually taken to indicate a distinction in the denotation: that is, both "the children" and "all the children" are taken to denote a group. Instead, it is argued that the contrast reflects a difference in how the predicate gets applied to this group. This proposal is summarized in the following hypothesis (from [Dowty, 1986b])

- (5.19) Hypothesis: the effect of all on a collective predicate is to fully distribute the predicate sub-entailments to every member of the group argument

Thus, (1.44b) results in a semantic interpretation that, although not identical with the interpretation of (5.15b), resembles it in that it does not introduce any scope dependency, hence the interpretation is that observed by Ioup

The treatment above is based on the assumption that "all" behaves in (1.44b) as a 'predeterminer' instead of as a quantifier. There are independent reasons for assuming that "all" is at least ambiguous between a quantificational reading and a modifier reading. For example, "all" may "float":

- (5.20) The children all got a present for their birthday.  
This ambiguity of "all" may explain the difference in scoping between "all the children" and "all of the children", displayed in (5.21):<sup>9</sup>

- (5.21) a John gave a box of candy to all the children  
b John gave a box of candy to all of the children

The preferred reading of (5.21a) is like the preferred reading of (1.44b): only one box of candy is involved. However, the preferred reading of (5.21b) is one that involves more than one box of candy. I propose that in (5.21b), "all" is interpreted as a quantifier.<sup>10</sup>

## 5.4 INDEFINITE NOUN PHRASES

### 5.4.1 Presuppositional and Non-Presuppositional Translations for Indefinite NPs

As discussed in §5.2, indefinites and other weak NPs such as many NPs have both a presuppositional and a non-presuppositional translation. The non-presuppositional translation for the indefinite "a boxcar" is as follows:

- (5.22) "a boxcar"  $\sim \lambda x P [\text{BOXCAR}(x) \wedge P(x)]$ , where  $x$  is a new discourse marker

This translation is immediately related to Heim's translation: the  $\lambda$ -indefinite NP introduces a discourse marker  $x$  and this discourse marker has to be new. The presuppositional translation for the same indefinite is as follows:

- (5.23) "a boxcar"  $\sim \lambda x P [x \models \text{BOXCAR}(x) \wedge P(x)]$ , where  $x$  is a new discourse marker

This second translation corresponds to Eng's 'specific' translation for indefinite NPs [Eng, 1991]. My proposal and Eng's are closely related. Eng generalizes Heim's theory by proposing of indefinites that each NP has two indices, instead of just one as in Heim's theory, both of these indices may be marked  $\pm\text{def}$ . The first index corresponds to Heim's file card index, an NP whose first index is  $+\text{def}$  must be familiar, i.e., its index must be identical with the index of some previously introduced NP. If the second index is  $+\text{def}$ , "a referent of the NP must be part of some set of objects already introduced in the context." A definite NP, according to Eng, is an NP both of whose indices are marked  $+\text{def}$ , a non-specific indefinite NP is an NP both of whose

<sup>9</sup>The example is due to Mich. Maron. [p.c.]

<sup>10</sup>The present treatment does not account for an often mentioned quantifier hierarchy effect, the average preference for each NPs to take wide scope than every NPs.

indices are  $-\text{def}$ , a specific indefinite is an NP whose first index is marked  $-\text{def}$ , whereas its second index is marked  $+\text{def}$ .<sup>11</sup>

In the proposal developed here, NPs, instead of having two indices, may have one or two parameters. As seen above, definite NPs have two parameters, indicating that both their referent and their resource situation have to be identified. Pronouns have only one parameter, for the referent. Where Eng would have an NP whose first index is marked  $+\text{def}$ , we have here a NP whose referent has to be identified in context and denotes a parameter.

The presuppositional reading of strong and weak quantified NPs, according to my theory, is the reading that corresponds to an interpretation in which the resource situation of the quantified NP has to be identified. Identifying the resource situation for the quantifier 'det P' clearly gives us a set of objects of type P in that situation; that is, where Eng would have an NP whose second index is  $+\text{def}$ , we have an NP with a parametric resource situation.

The difference between the two theories is that definiteness and specificity are taken to be syntactic characteristics of an NP according to Heim and Eng, whereas they are seen as pragmatic characteristics of an NP here. In some recent work on the definiteness effect [Zucchi, 1991; Ward and Burner, 1993] it has been argued that the definiteness of an NP is a pragmatic characteristics. I discuss in [Preson, 1994] a class of definite NPs that also seem to indicate that a treatment along pragmatic lines is more appropriate.

### 5.4.2 Logical Forms

Assuming that indefinite NPs are ambiguous results in the following prediction. The presuppositional interpretation of an indefinite NP is going to be treated as an operator, just like all other presuppositional interpretations of NPs; that is, the value of the logical form of an indefinite NP interpreted presuppositionally includes a sequence in which the indefinite is put in scope, thus may take wide scope over other operators in the sentence. The non-presuppositional interpretation, on the other hand, may only take narrow scope, as it does not have the possibility to raise. In other words, whether an indefinite NP may take wide scope over other operators in a sentence depends on whether it can be interpreted presuppositionally.

Since the choice between one or the other interpretation of an indefinite affects its scope, it is important for the theory proposed here to say something about the conditions under which one interpretation or the other is chosen. I propose the following tentative hypothesis: whether or not an indefinite NP is interpreted as presuppositional or non-presuppositional depends on whether it is interpreted as part of the 'given' part of information conveyed by the sentence or of its 'new' part [Clark and Haviland, 1977]. This definition may look circular: what I mean is that whatever rules are used to determine whether a sentence constituent is part of what's given and what's new in the sentence are going to determine whether an indefinite (and other weak NPs) is interpreted presuppositionally or non-presuppositionally. I will not discuss these rules here, just give some examples.

First of all, there seems to be a rule saying that the subject of a sentence typically corresponds to given information, whereas the 'predicate' ([Comp, TP]) typically provides new information to given information.

<sup>11</sup> Eng observes that the other combination cannot occur – if an NP is condefiniteness with some previously introduced NP, it must be included in some contextually specified set. She concludes that all definite NPs must be specific. I agree otherwise in [Preson 1994].

(This rule has been discussed in §1.5.) Thus, the indefinite "a tree" in (5.24a), that occurs in object position, and therefore is part of the new, is interpreted as non-presuppositional, whereas the indefinite "a kid" in (5.24b), that occurs in subject position, and therefore is taken to be part of the given, is interpreted as presuppositional.<sup>12</sup>

- (5.24) a Every kid climbed a tree  
b A kid climbed every tree

Pragmatic reasons may also play a role in determining whether an indefinite NP is interpreted as specific or not—for example, 'heavy' NPs such as "an old tree whose ancient roots came out from the ground"—that is, NPs that provide so much information about the type of object under discussion as to make it unlikely that the speaker is referring to an arbitrary instance of the type—typically will be interpreted as specific, so that the default specified above will be overridden in (5.25).

- (5.25) Every kid climbed an old tree whose ancient roots came out from the ground

My hypothesis about the rules that result in assigning the presuppositional or non-presuppositional interpretation to weak NPs are discussed in §6.9.

## 5.5 VERBS AND VERB PHRASES

The interpretation of verb predicates I adopt in this dissertation is, for the most part, that adopted in Montague Grammar, in which transitive verbs denote functions of type  $\langle e, \langle e, t \rangle \rangle$ , and intransitive verbs denote functions of type  $\langle e, t \rangle$ . As far as the syntax is concerned, I also adopt a very straightforward approach. The only matter that has to be discussed is the proposal that the translation of verbs includes information about their thematic roles. As discussed in §1.5, I suggest that the role assigned to an argument of a predicate affects the scope assigned to that argument. I discuss in Chapter 6 how this happens.

### 5.5.1 Thematic Roles

Argument selection—that is, mapping from grammatical relations into semantic arguments—is not a trivial matter. The sentences in (5.26), for example, show that the referent of a Subject NP can be mapped into different arguments of the predicate CLOSE.

- (5.26) a John closed the gate with the remote control.  
b The remote control closed the gate.  
c The gate closed.

The theory of argument selection I discuss in §6.3 depends on the notion of *thematic role* (thematic relation, 'theta-role,' 'case,' etc.) (Gruber 1965, Fillmore, 1968, Jackendoff 1972, Jackendoff 1983, Jackendoff, 1990, Chomsky, 1981, Carlson 1984, Chierchia, 1984, Wilkins

<sup>12</sup>There is a clear correspondence between this proposal and Dowty's Mapping Hypothesis, although according to Dowty the split occurs at I'.

1988, Dowty, 1989; Dowty, 1991]. This may raise some alarm, because thematic roles are one of the most abused notions in linguistics—it isn't clear if they are needed, and certainly there isn't much agreement as to what they really are. In the case of argument selection, however, a relatively elegant formalization can be achieved by using thematic roles. In addition, recent work of Dowty (1989; Dowty, 1991) and Jackendoff (Jackendoff, 1972; Jackendoff, 1987; Jackendoff, 1990), among others, has clarified the field considerably, and the approach to thematic roles developed there holds the promise to evolve into a concrete theory (i.e., one from which we can get actual predictions). I use here the 'minimalist' proposal about thematic roles made by Dowty, and in general make minimal assumptions about their use.

Dowty argues, first of all, that thematic roles as a semantic notion, to be formalized in model-theoretic terms.<sup>13</sup> Secondly, he notes that the only way for a theory of thematic roles to capture interesting linguistic generalizations is for the theory to deal with thematic roles types such as AGENT and THEME, as opposed to individual thematic roles such as KILLER or EATEN-OBJECT. The former are a way of expressing what certain arguments of different predicates such as *x* KILL *y* and *x* EAT *y* have in common: it's not clear what need there is for a notion like individual thematic roles (except perhaps for 'argument indexing').

Dowty defines in (Dowty, 1989) a *thematic role* (type) as "... a set of entailments of a group of predicates with respect to one of the arguments of each." For example, the subject arguments of the two-place predicates *x* MURDERS *y*, *x* DOMINATES *y*, and *x* INTERROGATES *y* all share the entailment that *x* does a volitional act, that *x* causes some event to take place involving *y*, etc. The role type AGENT is then defined semantically as "whatever entailments of verbs about NP referents are shared by the verbal argument positions that we label with the term 'Agent'."<sup>14</sup>

In (Dowty, 1991), a list of the traditional difficulties with the notion of thematic roles is presented. These difficulties include:

- Thematic role-type assignment is not always transparent. Notoriously difficult examples are sentences like "Nelson ran out of money" and "The circle surrounds the dot."
- 'Splitting' thematic roles more finely helps in resolving such difficulties in assigning a thematic role to a certain argument. For example, Jackendoff distinguishes between an AGENT and an ACTOR role: even more fine-grained distinctions have been proposed. How finely can we the various thematic roles be, before they lose their capability of expressing useful generalizations?<sup>15</sup>
- It's not always the case that two arguments may be assigned distinct roles. Predicates like *x* IS-SIMILAR-TO *y* or *x* WEIGHS-AS-MUCH-AS *y* are cases in point. Predicates that refer to commercial transactions such as BUY and SELL also originate difficulties of this kind: both buyer and seller must act voluntarily for such a transaction to take place.

<sup>13</sup>This position is not universally accepted. The  $\lambda$ - $\lambda$ -roles used in Government and Binding Theory (Chomsky 1981), for example, are a pragmatically synthetic notion.

<sup>14</sup>In Jackendoff's proposal as well, thematic role types are not treated as primitives, but are defined in terms of certain properties of the lexical 'verbs' used to translate predicates.

<sup>15</sup>It is perhaps interesting to note that much the same difficulty is encountered in dealing with two other popular notions in natural language processing, speech acts and 'thematic relations'.

The way out of these difficulties, according to Dowty, is to regard role types 'not (as) discrete categories' but rather as cluster concepts, like the prototypes of Risch and her followers (p. 571). In fact, Dowty argues that, as far as argument selection is concerned, we only need two types of these 'cluster concepts': *PROTO-AGENT* and *PROTO-PATIENT*. These two role types might be characterized by lists of entailments such as those given below, classifying a predicate argument according to its 'distance' from the prototypicals thus defined.

### Contributing properties for P-seem:

- a volitional involvement in the event or state  
b sentence (and/or perception)  
c causing an event or change of state in another participant  
d movement (relative to the position of another participant)

### Contributing properties for the P-patient:

- a undergoes change of state
- b incremental theme
- c causally affected by another participant
- d stationary relative to the movement of another participant

While most English transitive verbs have more than one such entailment for each argument ('should' for example has all of the contributing properties of P-agent for the subject, and all of the properties of P-patient for the object), in general these entailments are independent. For example, Dowty provides the following list of predicates that satisfy only one of the P-agent entailments for their subject

- a volition alone. *John is ignoring Mary.*  
b sentence/perception alone. *John knows/believes the statement. John sees/offers Mary*  
c causation alone. *Teenage unemployment causes delinquency.*  
d movement alone. *The rolling tumbleweed passed the rock.*

The traditional role types can be reformulated as combinations of prototypical entailments. Here is how Dowty proposes to define some of these role types:

**AGENT: volition + causation + sentience + movement**

**EXPERIENCER:** sentence without volition or causation

**INSTRUMENT:** causation + movement without volition or sentence

**THEME: change + incremental | theme + dependent-existence + causally affected**

**PATIENT:** as THEME, but without causally affected

The verbs whose subject satisfies the AGENT role thus defined and none of the entailments of p-patient, and whose object satisfies all of the entailments in the THEME role but none of those in the AGENT role, are called *primary transitive verbs* (see table 1). Most verb arguments are less prototypical.<sup>16</sup>

<sup>16</sup> Three of the common role types are missing from the list given by Dornqvist. They include SOURCE, GOAL, and BENEFACTIVE.

### 5.5.2 Semantic Transduction of Predicates

I propose to represent each contributing property to a proto-role type by means of a predicate of the form *PROTO-ROLE<sub>i</sub>*. Thus, I indicate the contributing property *a* to the cluster of properties of the proto-role AGENT ('volitional involvement in the event or state') as P-AGENT<sub>a</sub>, and the contributing property *c* to the cluster for the proto-role p-patient ('causally affected by another participant') as P-PATIENT<sub>c</sub>.

I use these predicates to characterize the semantic translation of a predicate as follows. I assume that the semantic translation of each predicate specifies, for all of the arguments of that predicate, which of the prototypical properties of either proto-*roll* that argument has. Take 'climb', for example. The first argument has a volitional involvement, is sentient, moves, and changes its state; the second argument is stationary relative to movement. This results in the following translation for 'climb':

$$\text{"climb"} \sim \lambda x \lambda y \text{ CLIMB}(x, y) \wedge \text{P-AGENT}_a(x) \wedge \text{P-AGENT}_d(x) \wedge \text{P-PATIENT}_s(y) \wedge \text{P-PATIENT}_v(y)$$

Providing such translations even just for the predicates found in the TRAINS conversations would be a task beyond the scope of this thesis, so I simply assume that each predicate has one.<sup>17</sup> In fact, for the most part I will not include the complete set of entailments when giving an example of predicate translation, for simplicity: the reader should 'however' keep in mind that translations that do not include the information about argument entailment are just shorthands

### 5.5.3 Syntax of Verb Phrases and Locative Forms

The s-structure of transitive verb phrases is shown in (5.28a), the position indicated as [Spec, VP] is typically empty, whereas [Comp, VP] is occupied by whatever complement the verb takes. The s-structure of intransitive verb phrases is shown in (5.28b).

(5.28) a. 
$$\frac{\text{VP}}{[\text{Spec, VP}] \quad \overline{\text{V}}} \quad \frac{\text{VP}}{[\text{Comp, VP}]}$$

Extensional transitive and intransitive verbs (such as Montague's "find") are not operators. The Cooper Value of a VP headed by such a verb is either a single sequence of the form {verb<sub>i</sub>} for an extensional transitive verb, where verb<sub>i</sub> is the interpretation of the verb; or else is obtained by generalized application of the interpretation of the head and the first element of each sequence in the Cooper Value of the complement. Generalized application is defined as follows

<sup>17</sup> Beth Levin's book (Levin, 1993) actually provides something that is very close to the kind of classification that is needed here.

**Definition 5.1** Let  $\alpha$  and  $\beta$  be CRT expressions. The generalized application of  $\alpha$  and  $\beta$  are:

$$\left\{ \begin{array}{l} \alpha(\beta) \text{ if } \beta \text{ denotes a function of type } (\tau, \tau') \text{ and } \alpha \text{ denotes an object of type } \tau' \\ \beta(\alpha) \text{ if } \beta \text{ denotes a function of type } (\tau, \tau') \text{ and } \alpha \text{ denotes an object of type } \tau \\ \text{undefined otherwise} \end{array} \right.$$

Intensional verbs such as "need" and verbs that take infinitives as a complement such as "want" are discussed below.

## 5.6 TENSE

Tense can obviously affect the interpretation of a sentence's constituents—for example, it can be the case the "John lived in Boston" without being true that John lives in Boston now—and contributes to the scopal ambiguity of a sentence. For example, the sentence "A student at this high school became a millionaire" may have two interpretations: one in which the person who becomes a millionaire is still a student at the high school, that can be represented by tense taking wide scope over "a student at his high school", and another one in which the person in question becomes a millionaire years after leaving the school, that corresponds to the interpretation in which the indefinite NP takes wide scope over tense.


### 5.6.1 The Syntax of Tense

This section and those that follow deal with aspects of phrase structure on whose proper treatment there is no agreement yet. I will therefore present the syntactic treatment I have adopted in some detail.

Both in the version of formal semantics developed by Montague (most prominently, in [Dowty 1979]) and in the generative tradition ([Stowell 1981, Chomsky 1980b, Haegeman 1991, Pollock 1989]) it is assumed that tense occupies a distinct syntactic position from the main predicate of a clause.<sup>14</sup> As discussed in § 3.3, I adopt here the analysis according to which tense syntactically is a functional category with its own maximal projection called TP. According to this analysis, TP corresponds to the category that is usually called S. This analysis has the advantage that it assigns a very clear syntactic position to tense operators, but otherwise nothing in what follows crucially depends on this particular syntactic analysis, and other analyses could be adopted, at the cost of making the construction rules more complex. An example of this analysis is the s-structure of (5.29) given in (5.30).

(5.29) John left.

<sup>14</sup> In GOSG, tense interval is introduced by the base rules [Gazdar et al. 1985].



In (5.30), the head of TP is occupied by the unrealized tense marker PAST. I assume that the contribution of tense to the semantics of a sentence is the translation of tense markers. I discuss the semantic translation of tense markers in the next section.

### 5.6.2 Contextual Factors in Tense Interpretation and the Semantics of Tense

There is a great amount of literature on the semantics of tense [Vendler 1967, Prior 1968, Bennett and Parée 1978, Dowty 1979, Kamp 1979, Bach 1981, Parée 1984, Eng 1981; Dowty 1986a, Webber 1987, Parsons 1990, Hwang and Schubert 1992, Kamp and Reyle 1993].

Authors up to, and including, Montague adopted Prior's treatment according to which the contribution of tense to a sentence consists of so-called tense operators such as P and F. Ever since Dowty's *Win d Meaning and Montague Grammar* [1979] his contribution has been most typically captured by assuming that tenses introduce existential quantifiers over times (or situations) (for example, that the sentence "Kim married Lee" means "there is a time in the past such that Kim married Lee at that time"). The translation of the tense according to Dowty [1979] is as shown in (5.31), where PAST(*t*) is true iff *t* is a time before the time of evaluation.<sup>15</sup>

$$(5.31) \quad (\exists t (PAST(t) \wedge AT(MARRY(L, J)))))$$

This characterization, however, only captures the most basic properties of tense. It has long been known that tense has 'anaphoric' properties—that is, that several discourse interpretation processes are involved in tense interpretation. Parée, in a famous paper [Parée 1973] noted that tense may have a deictic reading: the preferred reading of (5.32), for example, is the one in which the speaker is referring to a particular time in the past at which she didn't turn off the stove (as opposed to being a statement that never in the past did the speaker turn off the stove). Many authors (e.g., Eng [1981]) take this to indicate that tense should be treated *referentially*.

(5.32) I didn't turn off the stove. [Parée 1973]

There are obvious analogies between these cases of 'referential tense' and the examples of 'referential indefinites' discussed by, e.g., Fodor and Sag [1982]. Several authors (e.g., Kripke [1977]) argue convincingly that the referential properties of indefinites are part of the pragmatics of interpretation, rather than the semantics. The same argument can be made with respect to tense.

<sup>15</sup> In Dowty's treatment, tense is 'quantified' in a timeless sentence, and tense is introduced syntactically (Rule 5.39, page 330)—that is, there is no semantic translation for the syntactic PAST; the translation is added on by the semantic rule associated with the syntactic rule that generates a tensed sentence out of a tensed one.



The interpretation of tense in narratives is also subject to a phenomenon called *forward motion* (Hinrichs, 1981; Partee, 1984; Dowty, 1986a; Webber, 1987). If a sentence in a narrative contains an accomplishment or achievement predicate, that sentence is understood to describe an event occurring later than the time of the previous sentence's event. In (5.33), for example, the event of the president's walking over to John is understood to occur *after* the event of John entering the president's office.

(5.33) John entered the president's office. The president walked over to him. [Dowty, 1986a]

This phenomenon has been extensively analyzed in the DRT literature on tense. Kamp (1979), Hinrichs (1981) and Partee (1984) proposed versions of the DRS construction algorithm that captured forward motion by introducing into a discourse temporal relations between an event and the reference time introduced by the previous event in the narrative. This account suffers however from several problems, the most serious of which being that it has been noted (Dowty 1986a; Webber, 1987; Lascandes *et al.*, 1992) that not all cases of achievement/accomplishment following an achievement/accomplishment behave like (5.33).

(5.34) a John walked across Iowa. He started in Sioux City and headed east to First Dodge. [Webber, 1987]  
b John lost the race. His coach didn't prepare him well enough.

Dowty suggests that forward motion is a effect of a pragmatic default such as Grice's Ockler Maxim, rather than an aspect of the semantics of tenses.

Not only do events in narrative texts such as those just presented tend to 'move time forward' they also tend to be interpreted as part of the same 'thread' or 'story'. In part this is just what one would expect from Grice's maxim of relevance, there is more to this than that, however. For example, there are cases of *extended flashbacks*, such as (5.35) (similar to an example discussed in [Hwang and Schubert, 1992]).

- (5.35) a John and Mary went to buy a lawn mower  
b Somebody had stolen theirs the day before  
c They had seen the thief go away  
d John had run after him to no avail.  
e All the lawn mowers were too expensive.  
f They decided they couldn't afford a new one.

There are two narrative threads in (5.35), one describing John and Mary's visit to a store, the other John's chase of the thief. These threads correspond to the repeated use of the past and of the past perfect, respectively. The return from the past perfect in the past in also coincides with the return to the earlier thread. Note that (i) not only the simple past but also the past perfect can achieve an effect of continuity, (ii) more than one discourse reference time needs to be maintained in order to account for the return to an earlier one.

Neither of the properties just discussed is typical of narrative texts only, nor just of texts in the past. Consider, for example, the following fragment from the TRAINS dialogue in Fig. 1.2, 8.1.4. The sequence of events described by utterances 15.1-15.4 clearly constitutes a single thread of action, and the events in the thread are interpreted as consecutive, even though only once does the speaker use a clear indicator of temporal order ('and then' in 15.4).

UW	Speaker	Utterance
9.1	M	we should move the engine at Avon.
9.2		engine E.
9.3		to (inc)
10.1	S	engine E1
11.1	M	E1
12.1	S	okay
13.1	M	engine E1, to Bath, to (inc)
13.2		or, we could actually move it to Dunsville, and pick up the buscar there
14.1	S	okay
15.1	M	and look up the buscar to the engine.
15.2		move it from Dunsville to Corning.
15.3		load up some oranges into the buscar.
15.4		and then move it on to Bath

(5.36)

Hwang and Schubert (1992) developed a method for tense interpretation that addresses both the problem of finding which events in a text constitute a single thread and the problem of choosing the temporal relation between two consecutive events in the same thread. They propose that tense interpretation takes place in three steps. During semantic interpretation, tense is translated into the indexical tense operators *PAST* and *PRES*. For example, "John left" translates as in (5.37). Next, scoping takes place, the logical form of "John left" after scoping is shown in (5.38).

(5.37)  $[John < PAST > leave]$

(5.38)  $(PAST [John leave])$

The purpose of the third interpretation phase, called *deindexing*, is to assign to indexical tense operators an interpretation specific to the context of application. This is done with the help of data structures called *tense trees*. Tense trees are a way to compute the 'temporal parallelism' relations that are at play in examples like (5.35). After deindexing, "John left" gets assigned the interpretation in (5.39). An episodic discourse marker is introduced for each tense and aspect operator; the *\*\** operator discussed in §2.3 is used to specify the description of the episode; and the 'contextually loaded' predicate *ORIENTS* is used to relate the new episode to an episode in some 'thread'.  $e'$  is the last episode in some thread that also includes  $e$ ; the temporal relation between  $e$  and  $e'$  is to be determined by commonsense reasoning (say, along the lines of the proposal of Lascandes *et al.* (1992)). By default, *ORIENTS* is 'decontextualized' into a precedence relation if both  $e$  and  $e'$  are achievements or accomplishments.

(5.39)  $(\exists e \text{ BEFORE } now) \wedge (e' \text{ ORIENTS } e) \wedge ([John \text{ leave}] ** e')$

Tense trees are used to identify the episode  $e'$  that orients the new episode. Each node in a tense tree represents a certain level of 'temporal embedding'. The arcs of a tense tree are labeled with tense and aspect operators, an episodic discourse marker introduced while processing the temporal operator *O* that is in the scope of the temporal operators  $O_1 \dots O_n$  is stored away in the node of the tense tree *T* associated with the current narrative that can be reached from the root of *T* by following the path whose arcs are labeled  $O_1 \dots O_n$ . For example,  $e$  in (5.39) is stored in



at least at a different time, and almost certainly in a different town.<sup>20</sup> Finally, (5.45) has two possible interpretations: there was a single event of three students visiting the speaker (say, at 3pm), or there were three such events—say, a visit at 3pm, another at 5pm and a third one in the morning.

- (5.43) Every student entered the room  
 (5.44) All U.S. presidents were born in a small town.  
 (5.45) Three students came to see me today

## 5.7 ADVERBIALS

An *adverbial* is a prepositional phrase or sentential constituent of a clause that can be loosely characterized as specifying further properties of the eventuality described by the clause, as opposed to its predicate/argument structure that is specified by [Spec, TP] and the VP. Examples of adverbials are "at 3pm" in (5.46a) and "with a hammer" in (5.46b).

- (5.46) a The engine will arrive at 3pm  
 b John broke the window with a hammer

### 5.7.1 The Syntax of Adverbials

In the TRAINS conversations, adverbials occur in syntactic configurations such as

- (5.47) a [TP The engine will be in Dansville [pp at 3pm]]  
 b [TP The boxcar has been in Coming [pp for three hours]]  
 c [TP We should send an engine [pp to Bath] [TP to get the hoarcar]]

According to Reinhart (1983, p. 59–ff) there are two positions in which adverbials may occur: adjoined to VP, and adjoined to S (that is, adjoined to TP in the version of X-theory adopted here). Some adverbials are always verb-phrasal. Reinhart mentions instrumental PPs ("with a hammer") and manner PPs ("with care," "by car"). Adverbials that express subcategorized arguments are always verb-phrasal, too. Other adverbials are only sentential for example, causal complements with "although" or "if"; ally, there are adverbials that may be either verb-phrasal or sentential. For example, the locative PP "in his office" is generally interpreted as sentential in (5.48a), as verb-phrasal in (5.48b) (since the location is subcategorized for by 'place'), and again as verb-phrasal in (5.48c).

- (5.48) a Ben is an absolute dictator in his office  
 b Ben placed his new brass bed in his office  
 c Ben found a scratch in his office

Reinhart discusses some syntactic tests whose purpose is to classify an adverbial as either verb-phrasal or sentential. Examples of such tests are the pseudo-cleft tests. The first pseudo-cleft

<sup>20</sup> This should be compared with: 'All managers of ENI in the 50s were born in a small town'. ENI being the Italian Oil company. There was a time at which all managers of the company were from Maiera, a small Italian town.

test predicts that clauses of the form [TP [what ... did] is VP + PP] are ungrammatical for PP a sentential adverbial, acceptable for PP a verb-phrasal adverbial.

- (5.49) a \*What Ben is is an absolute dictator in his office.  
 b What Ben did is place his new brass bed in his office.  
 c What Ben did is find a scratch in his office.

The second test is complementary to the first and predicts that clauses of the form [TP [what did + PP] is ...] are ungrammatical for PP a verb-phrasal adverbial, acceptable for PP a sentential adverbial:

- (5.50) a What Ben is in his office is an absolute dictator  
 b \*What Ben did in his office is place his new brass bed  
 c \*What Ben did in his office is find a scratch

### 5.7.2 The Semantics of Adverbials

Hwang and Schubert (Hwang, 1992; Hwang and Schubert, 1995) propose a semantically based distinction between two classes of adverbials. Adverbials in the first class express properties of eventualities, and are thus called *episode adverbials*. Temporal and *c<sub>o</sub>-final* adverbials, for example, are translated as episode adverbials. Adverbials in the second class are translated as predicate modifiers, and in general express properties of *actions* (i.e. "eventual individual pairs"). These latter are called *action adverbials*. For example, in the sentence "John walked with Pluto in Disneyland," "with Pluto" is interpreted as an action adverbial and translates as a predicate modifier: "in Disneyland" is interpreted as an event adverbial and translates as a property of the walking event.

Hwang and Schubert's distinction between episode and action adverbials is orthogonal to the distinction between sentential and verb-phrasal adverbials noted by Reinhart. None of Reinhart's tests succeeds in distinguishing between these two classes of adverbials. The results of the pseudo-cleft test, for example, are as follows:

- (5.51) a What John did is walk with Pluto in Disneyland  
 (5.52) a What John did in Disneyland is walk with Pluto  
 b What John did with Pluto is walk in Disneyland.

Because syntactic tests do not differentiate among the two classes of non-subcategorized adverbials proposed by Hwang and Schubert, I propose to translate adverbials of both classes as predicate modifiers, but to differentiate them semantically. Roughly speaking, I propose to translate episode adverbials according to the following schema:

- (5.53)  $\lambda P \lambda x (R(x) \wedge P(x))$

in (5.53), R is the property of episodes that represents the contribution of the adverbial: this property is applied to the term 'x' denoting the current situation. Note that CRT expressions already denote (sets of) functions from situations to truth values; the translation above is thus

equivalent to Stump's system [1985] where the abstraction over situations is explicit. For example, "in Disneyland" translates as follows:<sup>21</sup>

$$\text{"in Disneyland"} \sim \lambda P \lambda x (INDISNEYLAND(x) \wedge P(x))$$


As for the class of adverbials that Harnag and Schubert call action adverbials, I propose to translate them according to the following schema

$$(S55) \quad \lambda P \lambda x (R(P)(x))$$

in (S55), R is the contribution of the adverbial, and this is directly applied to the predicate P. For example, "with Pluto," translates as

$$\text{"with Pluto"} \sim \lambda P \lambda x (WITH(PLUTO)(P)(x))$$

The sentence "John walked with Pluto in Disneyland" would thus receive the following truth conditions



### 5.7.3 Logical Forms

I take the position that adverbials are not operators, and therefore do not introduce scope ambiguities. This assumption may seem unwarranted in light of sentences like (S56), that can be interpreted either as stating that in many cases, when John smokes a cigar, it's 3pm, or that whenever it comes 3pm John smokes a great number of cigars. Nevertheless, there are reasons to believe that the two available interpretations of (S56) correspond in the two possible syntactic positions for the adverbial "at 3pm" — adjoined to TP, or adjoined to VP. For example, the two readings are only available when the adverbial may be both verb-phrasal and sentential adverbials that are only verb-phrasal are not ambiguous:

<sup>21</sup> Horrobin [1991] discusses a complication with temporal modifications also used by Levesley [1979] and Moore and Steedman [1998], that cannot be handled with this kind of translation. This problem is best illustrated by examples like the following

(S54) John left for 20 hours

Here the adverbial "for 20 hours" does not modify the eventuality of "leaving" — an other world it is not the leaving that lasts for 20 hours but the state of being away from the room. Horrobin notes that there is a whole class of verbs that behave like "leave". In order to represent this class of verbs, it would be necessary first of all to have an explicit representation of the result state as a part of the translation of a predicate, as well as to leave in the context to determine the eventuality which has the duration expressed by the "for" adverbial (i.e. make that eventuality a parameter of the translation).

<sup>22</sup> Horrobin [1991] makes a similar proposal

(S56) John smokes many cigars at 3pm.

When the adverbial is expressed by means of a PP, the complement of the preposition may be an operator: an example is the PP "in every town" in "Some reporters interview Kissinger in every town". In this case, what is put into store by the NP is "passed up".

Let's consider first of all the Cooper Value of PPs. This is defined as follows. Let prep be of the form

$$\lambda P \lambda P' \lambda x ((P \circ \text{PREP})(P'))(x)$$

(where PREP is the property contributed by the preposition) for an action adverbial, and of the form

$$\lambda P \lambda P' \lambda x ((P \circ \text{PREP})(x)) \wedge P(x)$$

for an episodic adverbial. We can then define the Cooper Value of the logical form of both PPs that express episodic adverbials and PPs that express action adverbials as follows

$$\bullet \text{ CV } \left( \begin{array}{c} \text{PP} \\ \diagup \quad \diagdown \\ P \quad \text{prep} \\ \diagdown \quad \diagup \\ \text{NP} \end{array} \right) = ((\text{prep} \circ \sigma_1, \sigma_2, \dots, \sigma_n)) \text{ for all } (\sigma_1, \sigma_2, \dots, \sigma_n) \text{ in CV(NP)}$$

The Cooper Value of the logical forms that represent VP-level adjunction of adverbials is defined as the "cross-product" of each sequence in CV(PP) and each sequence in CV(VP), each new sequence is obtained by applying the first element of a sequence in CV(PP) to the first element of a sequence in CV(VP), the new sequence also includes the rest of the elements of both sequences

$$\bullet \text{ CV } \left( \begin{array}{c} \text{VP} \\ \diagup \quad \diagdown \\ \text{VP} \quad \text{PP} \\ \diagdown \quad \diagup \\ \text{VP} \end{array} \right) = \left( (\sigma_1, \sigma_2, \dots, \sigma_n, \tau_1, \dots, \tau_m) \text{ for all } (\sigma_1, \sigma_2, \dots, \sigma_n) \text{ in CV(PP) and } (\tau_1, \tau_2, \dots, \tau_m) \text{ in CV(VP)} \right)$$

In defining the Cooper Value of the logical form obtained when sentence-level adverbials are considered (adverbials adjoined to TP), we have to take into account that the first constituents of sequences in the Cooper Value of TPs are objects of type  $t$ . In the case of adverbials that may only occur at sentence level, there is no problem: these can be defined as functions from sets of propositions to sets of propositions according to the schema

$$(S57) \quad \lambda \Phi / \Phi \wedge R(s'')$$

where, again, R is the contribution of the adverbial. The issue is what to do with those adverbials that may have either at VP level or a sentential level interpretation. Two routes are available either to assume a lexically ambiguity, or that one interpretation is more "basic", and the other interpretation is obtained by it. I follow the second route, and assume that the predicate modifier interpretation is more basic, an interpretation of the type in (S57) is obtained by "raising" the predicate level interpretation as follows

$$\lambda \Phi / \text{ADV}(\lambda x \ x = x\Phi^*) \wedge \Phi /$$

where ADV is the predicate modifier interpretation, the higher level is obtained by apply.  $\Phi$  this interpretation to the vacuous predicate  $\lambda x \ x = x$  and this in turn to the constant  $\Phi^*$  that is defined at every situation.

The Cooper Value of a logical form obtained by TP-adjunction of an adverbial consists of the set of sequences obtained by applying the interpretation of the adverbial (perhaps raised as above in case of ambiguous adverbials) to each first element of the sequences in the CV of the lower TP.

$$\bullet \text{ CV} \left( \begin{array}{c} \text{TP} \\ \text{TP} \end{array} \right) = \{ (\text{adv}^* \text{con}_1, \text{con}_2, \dots, \text{con}_n) \text{ for all } (\sigma_1, \sigma_n) \text{ in } \text{CV}(\text{TP}) \}$$

## 5.8 INFINITIVAL CLAUSES

I use the term *infinitival clauses* to indicate sentence constituents such as "to pick up the boxcar" that may occur, for example, in complement position (as in (5.58)) and in adjunct position (as in (5.59)).

(5.58) We need to send a boxcar to Corning.

(5.59) We need an engine to pick up the boxcar.

I review in this section the assumptions about their syntactic and semantic treatment adopted in the rest of the dissertation. (This material is assumed when discussing modals and adverbials, below.)

### 5.8.1 The Syntax of Infinitival Clauses

Two main positions have been adopted concerning the syntax of infinitivals: that they are VPs (e.g., in [Chierchia, 1984] and [Cazdar *et al.*, 1985]) or that they are sentential constituents (TPs, in the framework adopted here), e.g., [Haegeman, 1991].<sup>21</sup> Again, I follow the 'traditional' generative approach, and assign to the examples above the following s-structures:

(5.58') [TP We need [TP to [VP send a boxcar to Corning]]]

(5.59') [TP We need an engine [TP to [VP pick up the boxcar]]]

According to this analysis, the particle "to" is the head of a tenseless clause, the difference between the translation of infinitival clauses and the translation of tensed clauses comes from the difference between the translation of "to" and the translation of the tense markers discussed above.

<sup>21</sup> Jones [Jones, 1981] actually argues that infinitival clauses are to be treated as certain contexts as CPs, or other contexts as TPs, and in yet other contexts as VPs.

### 5.8.2 The Semantics of Infinitival Clauses

Infinitival clauses are translated into predicates in the PTO grammar, for example, "We need to talk" is translated as follows:


(5.60) We need to talk.

(5.60') NEED(we, TALK)

Chierchia however argues in his dissertation [Chierchia, 1984] that infinitival clauses denote *properties*—individuals in one-to-one correspondence with predicate intensions. Chierchia's position is adopted by Hwang and Schubert as well, who identify Chierchia's properties with Carlson's kinds [Carlson, 1978], more precisely, with *situation kinds*. The translation proposed by Hwang and Schubert for "We need to talk" is shown below.

(5.60'') NEED(we,  $K(\lambda s [TALK(we)] ** s))$

The translation I use is derived from Hwang and Schubert's, the \*\* operator is replaced by an event description, and the extra abstraction over situations is eliminated, as follows:



This translation is obtained by assigning to the particle "to" the following translation:

(5.61) "to"  $\sim \lambda \Phi. K(\Phi)$

Note that no events are introduced by this translation. I also assume that the empty category PRO that fills the position of [Spec, TP] translates in a parameter, and that the value of logical forms assigned to infinitival clauses is obtained from the semantic translation of the subject of such clauses, their complement, and the translation of the head much as the value of tensed TPs is (see above).

## 5.9 MODAL VERBS AND MODAL AUXILIARIES

### 5.9.1 The Syntax of Modal Auxiliaries


Trying to capture the syntactic properties of modal auxiliaries have proved very difficult, and there isn't much agreement in the literature. One of the problems is that modal auxiliaries appear in sentences with widely different syntactic structures. A sample of these constructions is presented in (5.62).

(5.62) a. We should send a boxcar to Corning.

- b We shouldn't send a boxcar to Corning.
- c The oranges should not be shipped in a boxcar.
- d The oranges should have been shipped in a boxcar.
- e The oranges should not have been shipped in a boxcar.
- f The oranges should have arrived in Corning by now.
- g Should we send a boxcar to Corning?
- h Shouldn't we send a boxcar to Corning?
- i Why should we send a boxcar to Corning?
- j Why shouldn't we send a boxcar to Corning?
- k Should the oranges be shipped in a boxcar?
- l Should the oranges have been shipped in a boxcar?
- m Should the oranges not have been shipped in a boxcar?
- n Should the oranges have arrived in Corning?


According to Chomsky's classical analysis in *Syntactic Structures* [Chomsky, 1957], modal auxiliaries, tense and negation are subordinate to a structure to a node called AUX that is a sibling of the topmost VP node. Tense, modal auxiliaries, and negation appear under the V node in the order Tense (Modal) have -en (be -ing).

#### Modal Auxiliaries in AUX




The AUX hypothesis is now generally thought to be unworkable, some of the arguments are presented in [McCawley, 1988]. In the versions of Government and Binding theory in which it is assumed that sentences are maximal projections of a functional category called Infl (e.g., [Chomsky, 1982, Chomsky, 1986a]), agreement, modal auxiliaries, and tense are supposedly located in the head of the maximal projection, i.e., so-called IP.

#### Auxiliaries in Infl: [Haegeman 1991]



With this kind of syntactic structure it is difficult to state generalizations about the restrictions on the order of verbs in the auxiliary complex one can see from the sentences in (5.62). An attempt at reconnecting the facts about auxiliaries with the restrictive form of X-theory developed by Stowell and Kayne [Stowell, 1981, Kayne, 1984] is made by Pollard [Pollard, 1989]. Pollard proposes to split IP in two separate projections, TP and AgrP, whose heads are Tense and Agreement, respectively. Optionally, a NEGIP maximal projection may occur in between. In Pollard's proposal, modal auxiliaries and tense are located under T, as follows:


#### Split Infl, Modal auxiliaries in T, and AgrP: [Pollard, 1989]



The tenseless components of the auxiliary complex go under AgrP, negation pives in NegP, eventual adverbs are adjoined in VP, or perhaps to AgrP.

Finally, there is the GPSG proposal [Gazdar *et al.*, 1985], according to which auxiliary verbs are heads of VPs. An example is shown below.<sup>24</sup>

#### Modal auxiliaries as heads of VPs: ([Gazdar *et al.*, 1985])




The theory of auxiliaries developed in GPSG and Pollard's proposal posit similar structural relations between modal auxiliaries and the VP. Also, NegP could simply be considered the same as the category AgrP+NEG of Gazdar *et al.* If we leave aside the issue of whether there is a category AgrP, the main differences between the two proposals, as far as this dissertation is concerned, are that (i) a functional category Tense is assumed in the GB literature, (ii) Tense

<sup>24</sup> For a similar analysis of auxiliaries in the government and binding framework, see [Zagona, 1988]. Zagona also proposes that auxiliaries are heads of VPs. McCawley proposes in *The Syntactic Phenomena of English* [1988] that the verbs are heads of V'. Another influential account of auxiliaries was proposed by Alimodian, Steele and Wexler [1989].

and modal auxiliaries fill the same position, [Head,TP]. The two theories may also differ in different predictions as far as the relative scope of quantifiers in subject and object position is concerned, however, it is assumed in some recent work on movement (Chomsky, 1986a) that maximal projections do not necessarily constitute barriers, so it's not implausible that the two approaches may end up making similar predictions as far as the scopability of operators are concerned, as well.

I propose a 'vanilla' variant of Pollock's proposal. I assume that modal auxiliaries are heads of TP, and that negation has its own projection, but I stay agnostic as far as the existence of an AgrP maximal projection between TP and VP. In this way the syntactic treatment I adopt could be replaced by something along the lines of Gazdar *et al.*'s proposal without affecting my theory of scope.<sup>27</sup> My analysis of "We shouldn't send a boxcar to Coming" is as follows:




The rest of the verbs in the auxiliary group are treated as heads of VPs, as in GPSG:

- (5.65) a The oranges should not have been shipped in a boxcar

<sup>27</sup>In the approach proposed by Pollock, it might be possible to explain the ungrammaticality of sentences like (5.61) by claiming that all constituents of the auxiliary group between the modal auxiliary and the past participle c in [Spec AgrP] thus ruling out the possibility that an adverbial may occupy a position between two kernels -components of the 'auxiliary block'. In an approach based on stacking VPs, and in which adverbs are treated as adjoining to VPs, this kind of sentence has to be ruled out by means of features. Given the currently poor state of knowledge about adverbs, however, this argument cannot be taken as decisive.

(5.63) \*The oranges should have usually been shipped in a boxcar



### 5.9.2 The Syntax of Modal Verbs

The syntactic distribution of modal verbs like "have to" and "need to" is illustrated in (5.66)

- (5.66) a We have to make OJ  
b We need to get a boxcar of oranges to Bath by 8am  
c We don't have to make OJ  
d We don't need to get a boxcar of oranges to Bath  
e We will/shall have to make OJ.  
f Do we need to make OJ?

While future auxiliaries are allowed with modal verbs, modification of modal verbs by other modal auxiliaries appears to be restricted, some of these restrictions appear to be pragmatic in nature

- (5.67) a We might/could/would have to make OJ  
b \*We should/can have to make OJ.  
c. We shouldn't/might not/wouldn't have to make OJ.  
d \*We can't/wouldn't have to make OJ.

There isn't much agreement on the syntax of modal verbs, either. The main topic of discussion is whether the infinitival complement of modal verbs should be treated as a VP (as proposed in GPSG or as a T'P-tense) (e.g., [Chomsky, 1981]). I adopt the latter solution, and assume the following structure for "We have to make OJ."

- (5.68) [TP [NP We] [vp have [TP [T-tense] to] [vp make OJ]]]

### 5.9.3 The Semantics of Modal Auxiliaries and Modal Verbs

Semantically, modal operators are usually treated as quantifiers over possible worlds. A typical clause for the truth definition of, say, the necessity operator □ will look something like (5.69) (from [van Benthem, 1988]), to be read as a model M satisfies the formula □φ iff the

possible world  $w$  iff  $M$  satisfies  $\Phi$  wrt  $z$ " worlds  $v$  that are accessible from  $w$  according to some relation  $R$ .

$$5.69) \quad M \models \Box \Diamond[w] \text{ iff for all } v \text{ with } R(w,v), M \models \Diamond[v]$$

it is also common to distinguish between different interpretations of modals. For example, the modal 'must' is called *deontic*, in the sense that it refers to a duty: in (570a) 'must' is called *epistemic*, because it refers to a piece of knowledge; in (570c) 'must' is called *dispositional*, because it refers to dispositions that people have. (Lyons 1977, Palmer 1986, Kratzer 1977)

5. 7(i) a All Maori children must learn the names of their ancestors.  
b The ancestors of the Maori must have arrived from Tahiti.  
c If you must sneeze, at least use your handkerchief.

As discussed in §1.5, Kratzer argues convincingly in [Kratzer, 1977, Kratzer, 1981] that those different interpretations of modal operators such as "must" can be accounted for without stipulating the existence of semantically distinct operators. Kratzer proposes a unique "relative 'must in view of'" operator as the common kernel of meaning for all the different interpretations, and argues that the different interpretations depend on how the contextual interpretation of modal base—*the modal base*—gets selected. The modal base restricts the worlds to be considered, the only worlds that have to be considered to evaluate the modal statement are those included in the modal base of its operator(s). For example, the sentences in (5.70) may be interpreted as follows ([Kratzer, 1977], p. 340).

- 571) a *In view of what their tribal duties are, the Maori children must learn the names of their ancestors*  
b *In view of what is known, the ancestors of the Maori must have arrived from Tahiti*

If this position is accepted, providing a semantics for modal auxiliaries (i.e. "must" and "should") requires a way to express the context-dependency of modalities. Chierchia and McConnell-Ginet (1990), for example, propose to make the interpretation of modals depend on an assignment of conversational backgrounds, in the same way that the context-dependency of pronouns is modeled by relativizing the interpretation to an assignment of values to use terms used to model translate pronouns (in their case, open variables). In (11)  $pr_i$  is used, if the propositions in the common ground at the *(worldtime)* pair  $(w, t)$  are  $\{p_i, p_j\}$ , the  $w, t$  is assigned by the variable assignment  $g$  at  $(w, t)$  is the set  $\{p_i\}$ . This set determines the modal force at  $(w, t)$ , which is the set of *(worldtime)* pairs  $(w', t')$  such that  $(w', t') \models p_i$  for all the propositions in  $g((w, t))$ . We use this system to define the semantics of "must" as follows:

The expression '*must*  $\Phi$ ' is true at  $\langle w, i \rangle$  relative to a background  $g(\langle w, i \rangle)$  iff the proposition that  $\Phi$  expresses follows from the propositions in  $g(\langle w, i \rangle)$  ( $\text{p} \subseteq_{\text{B}} \Phi$ ).

because my goal is to develop a theory about the pragmatic processes underlying discourse interpretation, processes which include determining the nodal base to be used to interpret any given utterance. I need a representation more explicit than Chierchia and McConnell Ginet's modalist analysis.

it should not come as a surprise by now that I propose to use parameters to model the context-dependent aspect of modality, much as I use them to model the context-dependency of quantifiers, definite descriptions and of tense

According to Kratzer's theory, a modal base consists of a set of circumstances, and it is identified by a set of propositions. In Situation Theory, and in particular in Epistodic Logic, a set of propositions identifies a *situation type*, and circumstances are replaced with real situations. As discussed in Chapter 2, one of the assumptions of the work in Situation Semantics is that the common ground does not constitute a uniform set of propositions, but is articulated into 'chunks' of information—situation types, in other words. In most work on Situation Semantics, only situation types that correspond to real world situations are considered (e.g., the visual situation or the discursive situation); it is but a small step, however, to allow for the common ground to include situation types that do not have a corresponding real situation. This move is additionally motivated, in conversations such as those in the TRAINS corpus, by the need to model reference to the plan (i.e. situation type) and the objects included in it; this issue is discussed more extensively in §6.1 below. Because situation types are in one-to-one correspondence with Huang and Schubert's situation kinds, the inclusion in the common ground of situation types without real world correspondence can be modelled in CRT by assuming that the common ground includes information about situation kinds.

Having made these assumptions, it becomes possible to characterize the task of selecting the modal base for a modal operator as the process of choosing one of the situation types/kinds in the common ground to provide the set of propositions that restrict the set of situations being quantified over by the modal. In other words, I propose that the epistemic and deontic readings of the statement "John must buy a car" can be paraphrased as in (5.72a) and (5.72b), respectively

- 572) a All situations  $s$  such that  $s$  is an instance of the discourse situation—the situation type consisting of the set of propositions that speaker and hearer share (see Chapter 4 and Chapter 6) are such that there is a situation  $s'$ , an extension of  $s$ , that include a sub-situation in which John buys a car
- b All situations  $s$  that are an instance of the situation type consisting of the set of propositions that characterize an obligation of the speaker,<sup>26</sup> are such that there is a situation  $s'$ , an extension of  $s$ , that include a sub-situation in which John buys a car.

Identifying the modal base for a modal operator means choosing one of these situations by pre-selecting situation kinds as the anchor for a parameter that is part of the translation of the modal. As in the case of quantifiers, the context-dependent elements are introduced by the translation of modal auxiliaries; the semantics of modals is thus defined independently from any particular context. I present the semantics of modal operators in detail in §4.3.<sup>27</sup>

<sup>28</sup> Note that, in general, one may question the assumption that all of the obligations of the speaker constitute a *maximal* set of obligations; that is, that it is *not* possible for there to be other obligations that do not allow for inconsistent abstract situations, as in (10). See, e.g., David Traum for this observation.

Graklen Katz (University of Illinois at Chicago) observes that the spectral class of stars decreases from O to M, and that the spectral class of galaxies increases from M to O. This is an interesting observation, but it is not a complete proof of a link between the spectral class of stars and the spectral class of galaxies. It is an interesting observation, but it is not a complete proof of a link between the spectral class of stars and the spectral class of galaxies. It is an interesting observation, but it is not a complete proof of a link between the spectral class of stars and the spectral class of galaxies.



An example of the translation assigned to modal auxiliaries is shown in (5.73). Modal auxiliaries are translated as functions from propositions to propositions, just like tense. According to this translation, the sentence "John must buy a car" is true in a situation  $s$  iff for all situations  $s'$  that are part of  $s$  and, furthermore, are instances of the kind of situation  $k$  to be identified, there is a situation  $s''$  that extends  $s'$  and includes a sub-situation that is an instance of the situation kind 'John buy a car'. The translation is shown in (5.74).

$$(5.73) \quad \text{'must'} \sim \lambda \phi. \text{must}([s' R k]) / ([s \subseteq s'] / [s' R k(\phi)])$$

$$(5.74) \quad \text{must}([s' R k]) / ([s \subseteq s'] / [s' R k(\text{buy}(a\text{-car}))])$$

Modal verbs such as "need" differ from modal auxiliaries in that their argument is a situation kind<sup>20</sup> instead of a predicate (as discussed in §5.8). I assume that infinitival clauses denote situation kinds<sup>21</sup>, otherwise the translation is identical to the translation of modal auxiliaries. The translation of modal "need" is as follows:

$$(5.75) \quad \text{'need'} \sim \lambda x_k. \lambda s. \text{MUST}([s' R k]) / ([s \subseteq s'] / [s' R k])$$

The result of the modal construction rules for modals are complex conditions of the following form:

$$(5.76) \quad [s' R k] \quad \text{must} \quad [s \subseteq s'] / [s' R k(\text{---})]$$

The denotations of these conditions is defined as specified in Chapter 4

<sup>20</sup> In an instance of "Thus if the situation in the nuclear scope is an instance of a telic predicate, the situation type is taken to be a type of course of action, and the parameter  $k$  will be expected to be a course of action that will result in a deontic interpretation, since obligations are typically obligations to act. On the other hand, if the situation in the nuclear scope is an instance of a stative predicate, and stative predicates are typically used to describe the background of situations, the parameter  $k$  will be interpreted as an instance of a situation type that consists of a set of propositions in the common ground.

<sup>21</sup> As discussed above, I follow here Chertus and Schubert in assuming that infinitival clauses denote situation kinds.


## 5.9.4 Logical Forms

A modal auxiliary can take scope over an NP in subject position, as shown by the example discussed in §1.5

(1.79) An engine must go to Avon.

This sentence has a reading in which "an engine" is in the scope of the modal auxiliary "must", and in order to get that reading we must treat "must" as an operator

I discussed in §5.6 how the Cooper Value of logical form fragments corresponding to TPs is defined. The translation of a sentence like (1.79) is again a TP, the translation of the modal auxiliary occupies the head position, much as the translation of tense does. We could do one of two things: either assume that modal auxiliaries are simultaneously a realization of a modal operator and of tense, thus their semantic translation already combines (5.73) with the translation of tense markers discussed in §5.6, or that two distinct operators occur in the position of [Head,TP]. The two positions are summarized in (5.77a) and (5.77b), where I have used the notation 'tense' to indicate the translation of tense, 'modal aux' to indicate the translation of a modal auxiliary, and 'tense + modal aux' to indicate a semantic object that combines the translation of tense and of a modal auxiliary.



The representation proposed by Hwang and Schubert is like (5.77b), in which two operators occur. This kind of representation would seem to lead to the prediction that two distinct scopes may be available: e.g., that (1.79) should have six distinct interpretations, some of which differ only in the relative scope of tense and of the modal operator. My intuitions on this matter are not very strong, but this prediction doesn't seem correct. On the other hand, leaving tense and modal operators separate results in a more modular theory, and also leaves open the possibility that additional operators may occur in the position of [Head,TP] in the case of more complex cases of modal auxiliary sequences, such as "must have been". So, I propose a compromise: modal auxiliaries and tense do have separate translations, but that these translations are combined when both of these operators occur in the same position in the underspecified representation, so that the logical form associated with a sentence with a modal auxiliary will be as in (5.78), where the translation of the tense is applied to the translation of the modal auxiliary.



other syntactic solution affects, however, the account of the interpretation process that results in assigning a scope to the operator in the PP. In fact, by treating the coda as a case of VP adjunction no further hypotheses are needed as discussed below. I am therefore going to adopt that analysis, but the reader should keep in mind that even in the case the NP adjunction analysis were chosen, there would be independently motivated principles that would result in assigning a scope to the operator in the PP.

### 5.10.2 The Semantics of There-Insertion Sentences and the Definiteness Restriction

The expletive "here" in subject position is generally assumed to be semantically vacuous.<sup>11</sup> As for the truth conditions of the clause as a whole, Barwise and Cooper (1981) propose that sentences like "There is a boxcar in every town" can be paraphrased as "A boxcar in every town exists." Keenan (1987) proposes instead that the sentence semantically is equivalent to the sentence "A boxcar is in every town." I follow Keenan's proposal, and I am going to attribute the difference in meaning between the two sentences to purely pragmatic factors.

The property of there-insertion clauses that has attracted the most attention is the so-called definiteness restriction. This is the fact, observed by Miltsark, that only certain NPs may occur in post-verbal position. Thus (5.87a) is acceptable, but (5.87b) isn't, at least according to most speakers.

- (5.87) a There is a dog in the garden  
b ??There is every dog in the garden

The Definiteness Restriction has been often formulated as a semantic constraint (Miltsark, 1974; Miltsark, 1977; Barwise and Cooper, 1981; Keenan, 1987). In recent work, however, this assumption has been challenged (Zucchi, 1993; Ward and Birner, 1993). First of all, it has been observed that there are exceptions to the restrictions on the occurrence of definites in post-copular positions. Counterexamples to the Definiteness Restriction that provide further support for its pragmatic origin are discussed in (Ward and Birner, 1993); one example is given in (5.88).

- (5.88) I'd love to get away from my job, the worried, the hills... I've thought of chucking it all and going to Hawaii. But there are the kids to consider.

On the basis of these counterexamples, it has been argued that the Definiteness Restriction is best explained as a restriction on the givenness status of the NP in post-copular position. This is in agreement with the observation that the most common pragmatic role of there-insertion sentences is to introduce new referents in the discourse.

Having redefined the notion of 'strong' and 'weak' NP in terms of parameters and existential presuppositions, we get a classification of NPs that is based on a pragmatic notion and at the same time gives us very similar predictions to Keenan's formulation on the basis of existential

<sup>11</sup>In the literature, it has been proposed that it serves as a locative indicator (Kuroi, 1971) and/or as a scope marker (Williams, 1984). This latter hypothesis is assumed under Heim's analysis below.

functions.<sup>12</sup> We can then formulate the Definiteness Restriction as a restriction on the occurrence of parametric NPs in post-copular position.

**Definiteness Restriction:** \*There be NP(x), where NP(x) is an NP whose denotation is parametric.

This definition allows for normally parametric NPs such as "the kids" to occur in post-verbal position whenever their parameter has been anchored (as in (5.88)).

### 5.10.3 The Definiteness Restriction and Scope

Heim, in (Heim, 1987), does not attempt to explain the Definiteness Restriction; she is concerned with deriving the restriction on the scope of the post-copular NPs observed above from the Definiteness Restriction. She observes that sentence elements semantically interpreted as variables, such as pronouns and traces left by wh-movement, cannot occur in post-copular position.

- (5.89) a \*Few people admitted that there had been them at the party  
b ??Which actor was there... in the room?

She proposes the following restriction to rule out cases like those in (5.89):

\*There be x, where x is an individual variable


If this restriction is assumed to hold at LF, i.e., if logical forms of the form *There be x* are ruled out by some semantic or pragmatic principle, then we have that the NPs in post-copular position have to be interpreted *in situ*, i.e., they cannot be 'moved out'.

This hypothesis is related to, but not identical with, the proposed redefinition of the Definiteness Restriction that I gave above, formulated as a prohibition against parametric objects in post-copular position. That restriction blocks pronouns, that are translated as parameters, but not variables.

However, if we take that definition, together with the suggestion in §5.4 that it is only the presuppositional interpretation of weak NPs that is subject to Quantifier Raising, we obtain the same effect that Heim intended to achieve with the formulation of the restriction as a restriction on the occurrence of variables. If only non-presuppositional interpretations of NPs are allowed in post-verbal position, and these interpretations do not behave as operators, then no wide scope reading can be obtained. In terms of the representation I have been using, this means that the (simplified) description in (5.91) is the maximally disambiguated representation of "There is a boxcar in every town" that may be obtained.

- (5.90) There is a boxcar in every town.

<sup>12</sup>Keenan claims that existential functions are the only objects allowed in post-copular position in there-insertion sentences. An existential function is a function from properties to properties such that for all properties  $p, q, p \in f(q)$  iff  $\exists x \in f(p \wedge q)$  where  $x$  is the vacuous property that is true of all individuals. I conjecture that the functions denoted by non-presuppositional interpretations of NPs are a subset of Keenan's existential functions.



The scope interaction between the pre-regular NP and other operators as the there-invention clause such as the modal 'must' in (5.79a) and the universal quantifier 'every town' in 'There is an equine in every town' is a direct consequence of Heym's constraint: the modal and the quantifier can be freely interpreted by the processes discussed so far, but the post-regular NP is forced to stay *in situ*, and therefore has to take a narrow scope.<sup>17</sup>

## 5.11 THE MEANING OF SENTENCES

To complete the specifications of the interpretation of logical forms we need to show how the final set of interpretations is put together out of the sequences  $\langle \alpha_1, \dots, \alpha_n \rangle$  obtained in (5.10a). This is the topic of this section.

### 5.11.1 The Cooper Value of CPs

5.11.1.1 The sequences in storage are displayed when computing the interpretation of CPs. This choice is motivated by the fact that CPs act as islands for scope: no quantifier can take scope outside a relative clause, an embedded clause with an explicit complementizer (such as 'what', 'x' 'whether'), or a sentence.

We are specifically interested in the Cooper Value of logical forms that are associated with CPs. As obtained from the Cooper Value of the embedded TP as follows:

$$\left\{ \langle \alpha_1, \dots, \alpha_n \rangle \mid \begin{array}{l} \langle \alpha_1, \dots, \alpha_n \rangle \text{ is obtained by taking a sequence} \\ \langle \alpha_1, \dots, \alpha_n \rangle \text{ in } \mathcal{C}(V(\text{TP})), \text{ generating a permutation of the operators in storage } \alpha_1, \dots, \alpha_n, \text{ all of which are functions from propositions to propositions, and then applying the first type-2 operator in the permutation to } \alpha_1, \text{ then the second operator to the result of the first application, and so forth.} \end{array} \right\}$$

<sup>17</sup> A counterexample to the generalization above is given in (5.92). This example is due to Janet Heyman.

(5.92) There is a town in Germany of precisely eleven buildings, all of which have eleven floors.

An interpretation is obtained by discharging all operators in storage in a sequence and by making the resulting proposition the denotation of a situation  $s$  to be determined. For example, consider the logical form for the sentence 'Every kid climbed a tree':

(5.93) Every kid climbed a tree

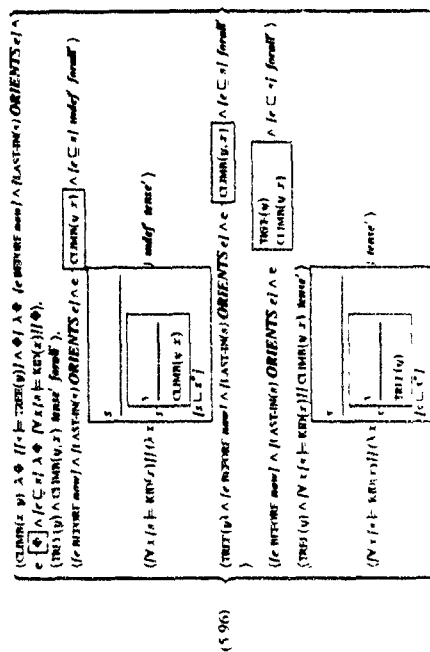


I have used *tense* as an abbreviation for the translation of tense. Starting from the bottom, the Cooper Value in the logical form fragment obtained from the VP is the following set of sequences:


$$\left\{ \langle \lambda x \text{ CLIMB}(u, x), \lambda \Phi / s \vdash \text{TRUE}(\Phi) / \lambda \Phi / s \rangle, \langle \lambda x \text{ TRUE}(u) / \lambda \text{ CLIMB}(u, x) \rangle, \langle \lambda x \text{ TRUE}(u) / \lambda \text{ CLIMB}(u, x) \rangle \right\}$$

Nine, that due to the ambiguity of weak NPs such as 'a tree' between a presuppositional and a non-presuppositional reading, three interpretations are obtained: only the presuppositional interpretation behaves like an operator, i.e., allows a situation prenotation, in what follows, all only consider two such interpretations: the first and the third.

Continuing bottom-up and applying the rule for assigning a Cooper Value to TPs, we obtain the following set (it has intensively used in all sequences except the first the abbreviations *tense* for the translation of tense, *index* for the translation of the indefinite, and *forall* for the translation of the universally quantifier  $\forall$  NP).



and so forth. The Cooper Value of the CP is then derived from (5.96) by means of the schema (5.97) and the denotation of the logical form is obtained from that schema yields a set of six readings for the sentence, two of which are listed in (5.97).




33. assuming that all operators in storage get 'discharged' when a CP is encountered, we only get the readings of a sentence that are consistent with the Scope Constraint. Quantifiers are blocked

by relative clauses and other embedded clauses with an explicit complementizer. On the other hand, quantifiers are not blocked by TPs. We can also explain the conflicting intuitions about the trace effect by assuming that the complement of *believe*, when no explicit complementizer is present, can be interpreted either as a TP or as a CP: depending on the interpretation, the clause works as a *sieve* or as an *island*.

### 5.112 A Revision


Some formally-inclined readers may have noticed that the schemas for assigning a Cooper Value to maximal projections, together with the definitions I gave concerning what value is put in storage for operators, do not return exactly the values listed in (5.96). The problem is that the system above does not include a way for binding discourse markers. It is possible to give definitions concerning which values are put in storage and how these are applied that guarantee that the markers are bound appropriately, but the resulting system is much more complex, so I have preferred to discuss more extensively the simpler system above. In this final part of the section, I will briefly discuss what changes should be made to obtain the desired translations.

First of all, we have to make sure that each discourse marker gets introduced as part of some substitution description from the very beginning. For example, we have to replace the definition below for the non-presuppositional translation of indefinite NPs with the definition that follows




598) "a boxcar  $\sim \lambda p \lambda s s$ "

A similar modification has to be made for all sentence constituents that introduce new discourse markers. Then, instead of putting in storage functions from propositions to propositions of the form shown in (5.96), we would have to put functions from propositions to functions from situations to propositions, so that the situation descriptions contributed by each sentence constituent can be merged.



599) "260KCM" ~  $\lambda \Phi \lambda s s$   $\wedge \Phi(s)$

**Tense gets (re)defined as follows**



(510) PAST  $\sim \lambda \Phi \lambda \lambda s \Phi(e$

The translation of predicates would have to be changed as follows, so that they introduce event descriptions (ignoring thematic roles)

$$(5.101) \quad \text{"climb"} \sim \lambda x \lambda y \lambda s \lambda t \left[ \text{CLIMB}(x, y) \right]$$

Finally, the translation for quantifiers should be modified so that the argument to the translation gets appropriately merged with the nuclear scope of the operator

$$(5.102) \quad \text{"every boxcar"} \sim \lambda \phi \lambda x \lambda s \left[ \text{every}(\lambda x \lambda s \left[ \text{BOXCAR}(x) \right]) (\lambda x \left[ \phi(x) \right]) \right]$$


These translations yield the desired result, although some complexity is involved in combining the interpretations

## 5.12 OTHER CONSTRUCTIONS AND SENTENCE CONSTITUENTS

I briefly discuss in this section some linguistic constructions and aspects of sentence meaning—negation, coordination, and apposition—that can be found in the TRAINS conversations, but whose syntactic properties, semantic properties, and scopal behavior I haven't yet thoroughly studied. The treatment I discuss appears adequate for dealing with the cases we have encountered, but should be seen as a provisional proposal

### 5.12.1 Negation

As discussed in §3.3, I assume that at  $\epsilon$  structure negation serves as the head of its own maximal projection, called NEG, optionally located between TP and VP



The semantic translation of negation occupies in a logical form the same position. The  $\epsilon$ -structure realization of negation is translated as a function from predicates to predicates, as follows

$$(5.104) \quad \text{"not"} \sim \lambda P \lambda x \lambda t \neg P(x)$$

As far as scope is concerned, negation currently is not treated as an operator: the Cooper Value of the logical form resulting from (5.103) is obtained from the Cooper Value of the complement VP merely by applying (5.104) to the first constituent of each choice, without putting anything in storage. This treatment has one main advantage, in common with the treatment of coordinating expressions such as conjunction: it results in plausible readings for the sentences we have looked at without making the assignment of a semantic scope to negation depend on discourse interpretation. This is an advantage because it's not at all clear what the presuppositional aspect of negation could be, although Kratzer makes some preliminary proposals in [1989]

One of the predictions of the current treatment is that a sentence like

$$(5.105) \quad \text{John didn't buy a car}$$

only gets the reading in which "a car" takes wide scope if the indefinite gets a presuppositional interpretation, that is, if it is interpreted as specific. This hypothesis is relatively plausible. A second consequence is that every operator whose position at  $\epsilon$ -structure is higher than the position of negation is predicted to take scope over negation. This is less satisfactory. For example, tense is always going to take scope over negation. (5.105) is interpreted as 'there is an event in the past that supports the proposition of John not buying a car'. This is indeed a reading of the sentence, however, there is another reading as well, that can be paraphrased as 'it is not the case that there is an event in the past of John buying a car'. Subjects are always going to take scope over negation as well, while this is arguably correct in the case of weak NPs (see [Horn, 1989]), it is known that in English negation can take scope over strong subjects, as in the following example


$$(5.106) \quad \text{All boxcars didn't arrive in time}$$

This sentence has a reading that can be paraphrased as: 'It is not the case that all boxcars arrived in time'

The simplest modification to the current treatment  $\rightarrow \epsilon$ , one that allows negation to take wide scope over tense and subjects, but falls short of treating negation as an operator—is to assume that negation is located in the head of TP. Given the way that the Cooper Value of TPs is defined in §5.6, the required readings would be produced. The assumption that negation occurs in [Head, TP] is not as well motivated as the position that modal auxiliaries occupy (that is, 'I don', however, for example, sentences such as "John did not buy a car" seem to call for a treatment in which tense and negation occupy distinct positions. There is a problem for scope as well, namely explaining why when tense and negation both occur in the position of [Head, TP] they can take either relative scope, whereas when a modal auxiliary and negation occur in that position only one scope is available (that in which tense takes wide scope over the modal auxiliary))

### 5.12.2 Coordination

I assume the generalized treatment of coordination proposed in [Partee and Rooth, 1983], where conjunction and disjunction are treated as polymorphic operators that can coordinate NPs and VPs as well as CPs. Neither conjunction nor disjunction are treated as operators, the Cooper Value of coordinated structures such as (5.107) is obtained by applying the translation of the coordinator to the coordinated structures



I assume that *serial coordination* is a coordination of CPs; this makes it a scope island in the sense that the operators occurring in either coordinated CP get 'discharged' (i.e. taken out of storage) by the *it* that specifies the Cooper Value of the coordinated CPs; thus cannot take scope over operators occurring in the other coordinated CP.<sup>14</sup>

### 5.12.3 Apposition

Appositions and non-restrictive relative clauses are extremely common in the TRAINS dialogues.

(5.108) Move the engine at Avon, engine E1, to Bath to pick up the boxcar

(5.109) Move the engine to Avon, where there are oranges

We currently treat both appositions and non-restrictive relative clauses as cases of adjunction but much more needs to be said. McCawley (1988) takes these constructions as arguments for the existence of 'crisscross' syntactic trees. As far as the pragmatics of discourse is concerned, these constructions introduce speech acts that are subordinate to the main speech act performed with an utterance. In (5.109), for example, the clause 'where there are oranges' provides an explanation for the decision of moving the engine to Avon. We are elaborating a treatment of these constructions based on the idea of 'micro conversational events' discussed in §3.8.6.

<sup>14</sup> A weaker form of Scope Constraint can be obtained by allowing sequential coordinations of TPs as well. Since no discharging takes place at the TP level the operators in one of the coordinated TPs could take scope over the operators occurring in the other TP.

## 5.13 SUMMARY OF THE CHAPTER

In this chapter I spelled out in detail my syntactic and semantic assumptions, and discussed the effects of these assumptions on semantic ambiguity and scope.

First of all, I specified which sentence constituents behave as operators, and which scope they can take. I related the distinction between operators and non-operators (sentence constituents subject to Quantifier Raising and sentence constituents that are not, according to the terminology adopted in generative grammar) to the distinction between presuppositional and non-presuppositional objects. I also discussed how syntactic constraints can be enforced by means of underspecified representations.

Most importantly, I claimed that the sentence constituents whose interpretation needs to be completed in context (i.e., the sentence constituents with a parametric interpretation) are exactly those constituents that have been called in the literature 'presuppositional', such as strong NPs, and for each of these constituents I indicated what part of their interpretation is parametric. This classification is an important preliminary to the discussion in Chapter 6, in which I discuss the discourse interpretation procedures that assign a value to these context-dependent elements.

## 6 Discourse Interpretation and Scope Disambiguation

The notion of scope is a way to capture **informational dependency** relations between the interpretation of sentence constituents. A sentence constituent depends on some other sentence constituent, or on some object in the common ground, if its interpretation is affected by the interpretation of this sentence constituent or object in the common ground.

Throughout this dissertation, I have often used the term 'assign a scope' between scope quotes. This is because I think that the formulation of the Condition on Scope Disambiguation I gave in Chapter 1 should actually be strengthened as follows:

**Condition on Scope Disambiguation (Strong Version)** A listener processing a sentence is not concerned with assigning a 'scope' in the semantic sense to sentence constituents, but with establishing informational dependency relations

What this means is that the reasoning that results in a particular scope ordering consists for part of inferring anaphoric reference relations, resolving presuppositions relations, or partitioning a sentence between what is given and what is new. While the outcome of these processes is an interpretation that can be formulated in terms of a formal semantic theory such as CRT, the process by which this interpretation is arrived at has nothing to do with 'scope' per se. Once we factor out the effect of syntactic and semantic constraints, the scope preferences observed in the literature on scope interpretation and those discussed in §1.4 can be explained as the result of the interaction of several discourse interpretation processes such as definite description interpretation and the interpretation of modals, that operate on the basis of the information provided by underspecified representations.

If we think of scope assignment in this way, rather than a way for establishing structural relations in logical forms, we get a different way of looking at the process of assigning a scope to operators. Instead of thinking of it as a process in which listeners are concerned with choosing one among many several permutations of operators, we can look at it as a process during which listeners try to establish such informational dependency relations, and arriving at an interpretation that can be specified in terms of 'traditional' scope relations during the process.

In this chapter I elaborate on, and present support for, the thesis above. I analyze in detail two discourse interpretation procedures that occur in all kinds of conversations, the compilation of the given/new partition of a sentence and argument selection, as well as the processes involved in the interpretation of the most common operators found in the TRAINS dialogues, and I

show how scoping preferences arise from the interaction of these procedures. The treatment of syntactic and semantic information introduced in Chapter 5 is assumed, in particular, I will use the proposals concerning which sentence constituents are presuppositional, and which aspects of their interpretation needs to be filled by contextual reasoning.



## 6.1 THE DISCOURSE SITUATION IN THE TRAINS CONVERSATIONS

Before discussing discourse interpretation it is necessary to look in some detail at the mutual information that is used and modified by the discourse interpretation procedures. In Chapter 3 I argued for a theory of the common ground oriented towards a representation of the discourse situation, as opposed to the topic of the conversation, and I showed how Conversation Representation Theory can be used to represent aspects of the discourse situation such as discourse topics and the focus of attention. In this section, I look more in detail at one case of discourse situation, the discourse structure in the TRAINS conversations, and how I propose to represent it in CRT.<sup>1</sup>

As discussed in §3.7, a typical TRAINS conversation is 'about' more than one discourse topic. The common ground includes

- 1 facts about the world: the world, in the TRAINS conversations, can be identified with the world described by the map
- 2 generic information about the task, such as expectations about the intentions of each conversational participant, and causal information (e.g., after unloading a boxcar, that boxcar becomes empty)
- 3 information about what has been said (the 'discourse history')
- 4 the current status of the plan

There are some interesting complications involved in modeling the kind of discourse topics that occur in the TRAINS conversations, especially at the light of the proposal made in §3.8.4 that a discourse topic is really a situation, namely, the situation that a conversational thread is describing. First of all, as we have seen in §1.4, it is possible for the speaker to use the plan (that is not, on the face of it, a 'situation') as a discourse topic and to refer anaphorically to objects in the plan, even when these objects are not (yet) related to any specific object in the map. As an example, consider the fragment in (6.1): the discourse marker 'a boxcar' is subordinate to the modal 'need,' and is interpreted as referring to an object in the plan, yet it may serve as an antecedent to the definite NP 'the boxcar'.<sup>2</sup> One problem to be addressed, then, is how to represent the plan, and in which sense can the plan serve as a discourse topic

(6.1) U so, we need a boxcar to move the oranges

S right

U okay, so we need an engine to move the boxcar

A second problem is that some of these 'discourse topics' are defined *intensionally*: i.e., the agents are aware of their existence even though they may only have a partial characterisation of

<sup>1</sup>For simplicity I am only concerned here with those situations and situation types which appear to be part of the common ground during the whole conversation. Judging from our conversations, it appears that the participants are able to create additional shared private situations dynamically. At the moment, it is not clear to me how this process works.

these situations. This is the case with the visual situation, for example: the participants to the conversations are clearly aware of its existence, but do not have a complete representation of it at any time.<sup>3</sup> I propose below that for each intensionally defined situation there is a situation forming principle which states the conditions under which a conversational participant can assume that a certain piece of information is part of that situation. These situation forming principles are mutually known to the participants, and can be used to assume that a certain fact is mutually known.

It may be useful before proceeding to remind the reader that the indexical *onstant ds* refers to the discourse situation, whereas *s* is an indexical that refers to the current situation.

### 6.1.1 (Visual) Information about the Map

One of most important discourse topics in the TRAINS conversations is the 'visual situation', that, in our case, is the world represented on the map. The participants to the conversations treat information about the map as if it were information about an actual situation. I call this situation *map situation* (MapS). MapS is the resource situation for definite description vis like 'the boxcar at Elmira' (1.3.3 in (1.61), §1.4) or 'the tanker car' (1.3.8), that are interpreted with respect to the 'visual situation'. The information in MapS represents the 'visual field' of the agents.

Although MapS is an actual situation, it is defined intensionally, in the sense that not all information contained in the map need to be part of the common ground at all times for the conversational participants to be able to explain that information in definite descriptions. That is, an agent may refer to objects contained in the map situation without knowing whether the other agent is aware of the existence of that object; it's enough for the speaker to assume mutual knowledge of the existence of that situation.<sup>4</sup> The TRAINS conversation, for example, include exchanges like:

- (6.4) 1 A You see the boxcar at Danville?  
2 B wait, got it.

The following situation-forming principle is associated to the map situation:

**Situation Forming Principle 6.1.1.1** *The conversational participants can assume that the fact that a proposition  $\Phi$  is supported by MapS is mutual knowledge if the map that the conversational participants are looking at serves as the source of information for that proposition.*

<sup>2</sup>This situation is confirmed by recent experiments by Mary Hayhoe at the University of Rochester, from which it appears that participants to visual tasks scan the visual scene continuously—i.e., they do not assume the information from the 'visual scene' is a constant set of propositions—and have difficulty performing their task if this scan is made hard or impossible.

<sup>3</sup>This is a general property of the visual field. Webster, for example, discusses examples like the following:

(6.2) Go to the kitchen, get THE COFFEE POT, and come back.

(6.3) Open the box, take THE RED BALL, and give it to me.

In neither example the listener is required to have prior knowledge of the existence of the object referred to by the speaker in order for the definite description to be felicitous.

This principle allows the participants to use the information in the map as if it were part of the common ground at all times.<sup>4</sup> I assume from now on that most DRS contains a discourse marker *Maps* referring to the map situation, and that facts in the map can be used even without previous mention, so that, for example, a participant can refer to 'the buscar at Avon' without having mentioned it before.

### 6.1.2 The information about the Plan

Of the discourse topics in the TRANS conversations the most important is of course the plan. The conversations consist in large part of statements like (6.6) that are interpreted as contributions towards the development of the plan.

(6.6) We should send an engine to Avon.

That the plan represents a discourse topic in its own right is also shown by the fact that the participants in the TRANS conversations often refer anaphorically to objects that have only been mentioned in statements like (6.6); this is what has not been connected with an object in the world. Thus (6.6) could be followed by

(6.7) We should hook up the engine to a boiler

How do these statements fit with the hypotheses about discourse topics and reference discussed below? In which sense can we say that the plan serves as the discourse topic of a statement like (6.7)? (As we will see later, this issue is related to the problem of representing modal subordination.) And how can we relate this fact with the location theory of definite descriptions interpreted discussed in § 3 (Hawkins, 1978; Clark and Marshall, 1981), according to which interpreting definite descriptions requires identifying the situation of which the object referred to is a part, and that resulted in the assumption that discourse topics have been said to be situations? The plan is certainly not an actual situation, and one may wonder whether a plan should be considered a situation at all. For example, the hypothesis about the ontological status of plans currently prevailing in Artificial Intelligence is that plans are 'recipes' (to perform kinds of actions, a recipe being a graph whose nodes represent 'operations' (action types) and whose arcs represent temporal or causal relations (Fikes and Nilsson, 1971; Sacerdoti, 1977; Allen *et al.*, 1990)).

Although the steps of the plan can be specified both using modal statements and using imperatives, imperatives are most common, and in general may be argued to be the prototypical

<sup>4</sup> This situation-forming principle is a corollary of Clark and Marshall's 'physical experience' heuristic (Clark and Marshall, 1981).

<sup>5</sup> Some speakers appear to assume the existence of two map situations. These speakers appear to assume the existence of a shared map situation in which no *Maps* has occurred, and of a second map situation in which the *Maps* has occurred. Only this second map is actually shared. Evidence for this is the abundance in our dialogues of utterances like (6.5) whose purpose, at least in part, is to make sure that certain information is actually shared. (These acts may also be interpreted as indirect requests for clarification (Gibson 1984) and indirect ways to impose constraints on the plan (Traum, 1991)).

(6.5) U: There is a furnace at Avon

form for expressing instructions. Because plans can be seen as complex instructions, it seems reasonable to have plans denote the same type of objects as imperative sentences. Having said Scholten assume that imperatives denote situation kinds (p.c.). I adopted their treatment in [Friesen, 1993], and used situation kinds to model plans as well. More precisely, I propose there to introduce special objects called *recipes-for-action*, a particular type of situation kinds, to model plans, among other properties, recipes-for-action have the property of being kinds of courses of action, as opposed to kinds of arbitrary situations. I use below the term *Plan*<sup>6</sup> to indicate the plan, and assume that the plan is a recipe-for-action.

I also propose in the paper that the specification of a plan does involve the construction of a *possible situation*, what I call *plan situation*, that is used as the resource situation of definite descriptions referring to objects mentioned as part of the plan. (A possible situation consists of information about other possible worlds—for example, worlds in which the events which are part of the plan actually occurred.) The plan itself is obtained by abstracting over this situation.<sup>7</sup> The analysis of modal subordination adopted here dispenses with the plan situation and is instead more directly related to Roberts' approach to modal subordination, as augmented in [Friesen and Zucchi, 1992]. I discuss next the processes that lead from modal statements to augmenting a plan, and my proposal concerning modal subordination, in the section on modals below.

In what follows, both the existence of a *Plan*<sup>8</sup> and the fact that the task of the user and the system is to develop a recipe for the kind of action they have to accomplish, are assumed to be part of the common ground.<sup>9</sup>

### 6.1.3 Discourse Topics: Situations vs. Situation Types

Although I have tentatively concluded that both situations and situation types/kinds may serve as possible discourse topics, one may wonder whether it would not be simpler to assume that it's always situation kinds that serve as discourse topics. For example, instead of assuming that when speaker and listener are talking about the state of the world, their discourse topic is a particular situation called *Maps* as I have proposed above, we might say that the discourse topic is a situation type consisting of the set of propositions about the map.

This approach is more uniform, and has some appeal when considering what might be the discourse topic of conversations that deal with some facts about the world (say, some property of boaters). However, it seems more intuitive to me to assume that when speaker and listener are talking about the world in which they are operating, the world itself serves as the discourse topic, rather than the set of propositions that are true of the world. This seems especially important when we consider the 'intensional' nature of some of these discourse topics discussed above: that is, the fact that even though the participants to a conversation may be talking about

<sup>6</sup> Note that when Gibson and Scholten propose in [Gibson and Scholten, 1986] that all objects and events introduced in the common ground are put into a 'focus space stack' used to interpret referring expressions, they are implicitly assuming the existence of such a possible situation.

<sup>7</sup> I nevertheless wonder whether there is a very psychological validity in the idea that an instance of the plan is actually created when developing it.

<sup>8</sup> Again, this might be an application of Clark and Marshall's 'linguistic competence' heuristic: since everything which is part of *Plan*<sup>8</sup> is introduced in the discourse




Figure 6.1 The common ground in the TRAINS conversations—a summary

the world they need not be aware of all that is true at the world. This seems to be suggesting that it is the world itself that serves as topic of discussion

#### 6.1.4 Summary

The structure that I propose for the common ground in the TRAINS situations is summarized in Fig. 6.1

## 6.2 THE STATUS OF INFORMATION IN DISCOURSE, THE GIVEN/NEW ARTICULATION, AND SCOPE DISAMBIGUATION

So far I have mostly been concerned with the representation of dependencies between an utterance and the common ground, such as reference relations and relations of presupposition. In this chapter, as well, I will mostly be looking at the processes that establish these kind of relations. There are also processes that establish dependency relations between the constituents of the same sentence, however; I discuss one in this section.

If we think of scope as a way to encode informational dependency, we would expect scope preferences to be affected by the status of the information communicated by an utterance. It should matter, for example, whether an NP conveys 'given' or 'new' information: the 'new' information depends on the 'given' information for its interpretation, so we would expect that the 'new' information will tend to take 'narrow scope' with respect to the old.

Prime factor, this hypothesis is consistent with what is known about the relative scope of NPs and other given/new status indefinites, that typically express new information, tend to take narrow scope, as seen above, one might argue that it is only specific indefinites that may take wide scope, and these have a 'given' component. Definites, on the other hand, that are typically interpreted as part of the 'given' information, tend to take wide scope, except when interpreted associatively or anaphorically to other operators.

In this section I am going to look in more detail at what has been called the given/new contrast by Clark and Haviland—the use by the speaker of an utterance of syntactic information to inform the listener of what's given and what's new—and the relation of the contrast to scope. I expand on this section on the argument discussed in §1.5 that both the preference for subject NPs to take wide scope and the preference for NPs in fronted position to take wide scope are the result of the way the given/new partition is computed.

### 6.2.1 The Given/New Contrast

Clark and Haviland (1977) argue that if we really believe in one of Grice's Conversational Maxims, the Maxim of Relevance (Grice, 1967), we are led to conclude that

. . . To ensure reasonably efficient conversation, the speaker and the listener adhere to a convention . . . the speaker . . . agrees to convey information he thinks the listener already knows as given information, and to convey information he thinks the listener doesn't know as new information.

Clark and Haviland agree that the distinction between given and new information is inherently coded in certain syntactic constructions, and that this coding can be revealed by linguistic tests. For example, a question such as (6.8), that conveys information about what the speaker already knows as well as his/her explicit indication of what he/she doesn't know, can be felicitously followed by one of the a sentences in (6.9) and (6.10), but not by one of the b sentences

(6.8) Who broke the window?

- (6.9) a It was John who broke the window  
 b #It is the window that John broke
- (6.10) a JOHN broke the window  
 b John broke THE WINDOW.

Clark and Haviland's interpretation of these data is as follows. Both it-cleft sentences and stress are ways of coding the given/new partition in a sentence. "It was N that G . N expresses new information, whereas G expresses given information, thus (6.9a) is a felicitous continuation of (6.8), as it provides the required new information, but (6.9b) isn't, since it amounts to a claim that the information that the window was the theme of the breaking event is new, which it isn't. Similarly in (6.10) a stressed sentence constituent conveys new information, thus (6.10a) is felicitous, but the b sentence isn't.

Clark and Haviland also note that whereas certain linguistic constructions are used with the express purpose of conveying the partition, others only do so more weakly, thus for example subjects are by default taken to express given information, but this preference may be overridden (e.g. by using stress, as in (6.10a)).

## 6.2.2 Formalizing the Given/New Contract

I assume the following about the correlation between syntactic structure and the given/new partition

- It-clefts, stress, and PP preposing<sup>10</sup> are all 'strong' methods to indicate the given/new partition
- The subject/predicate distinction is a weak indicator of the given/new partition

I further assume that 'givenness' is very similar to presupposition as discussed in §3.6, and 'newness' is closely related to 'novelty' in the sense of Heim – a structural indication that a certain sentence constituent is 'given' translates, in terms of operations on the common ground, into a request to either relate the sentence constituent in question to some element already existing in the context or else add such an element to the context if not already there.


The following rule formalizes in GWT the preference for subjects to be interpreted as providing given information. In the rule in (6.12) it is assumed that what's given is the property on which the quantified NP lives, the rule states that if a strong NP occurs in the position of (Spec,TP), a situation description can be added to the context, and the situational parameter representing the resource situation of the quantified NP can be activated in that situation description. This rule

<sup>10</sup> The notion of PP preposing is motivated by examples such as

- (6.11) a Where did some reporter interview Kravages?  
 b Some reporter interviewed Kravages in every town  
 c #In every town some reporter interviewed Kravages

That preposing is a strong indicator of 'givenness' is also often stated in the functional literature, e.g. in Kravages (1981). An extensive analysis of preposing can be found in Ward (1991).

can be interpreted as a rule for accommodating into the context the presupposition associated with a strong NP in case that setting NPs in subject position is



I emphasize that SUBJECT GIVEN applies if the quantifier in subject position is presuppositional, i.e. if the NP in subject position is either strong, or if the presuppositional interpretation of a weak NP has been chosen. A rule that differs from SUBJECT GIVEN only in its trigger formalizes the preference for interpreting NPs embedded in preposed PPs as given, this rule let's call it PREPOSED-PP GIVEN, is a strong default instead of a weak one.

## 6.2.3 Givenness and Scope Preferences

It is easy to see how the rule presented above can account for the preferences displayed in sentences like (6.13).

- (6.13) Every kid climbed a tree.

Unless some discourse interpretative process with higher priority has already determined the resource situation of the quantified NP "every kid", the rule SUBJECT GIVEN is activated. Once the resource situation is activated, the model construction rules are enabled, this results in assigning to "every kid" wide scope over "a tree". (The results of inference in this example have been presented in §1.5.)

The contrast observed by Heimart between (6.14a) and (6.14b) can also be accounted for. The fact that the reading in which "every town" takes wide scope is much more preferred in

<sup>11</sup> The actual proposal is to add a set of objects of the appropriate type to the context. An alternative implementation in terms of kinds (which also be formalized)

(6 14a) than it is in (6 14b) can be derived from the fact that the rule PREPOSED-PP-GIVEN is a strong default, whereas the rule SUBJECT-GIVEN is a weak default

- (6 14) a In every town, some reporter interviews Kissinger  
b Some reporter interviews Kissinger in every town

### 6.2.4 The Mapping Hypothesis

Diesing's *Mapping Hypothesis* [Diesing, 1992] is also a proposal about the the mapping from a syntactic representation into the kind of representations assumed in DRT, i.e. about how the material in a quantified sentence is divided between the restriction and the nuclear scope<sup>11</sup>

*Mapping Hypothesis:* Material from the VP is mapped into the nuclear scope  
Material outside the VP is mapped into the restriction

Intuitively, this way of splitting the sentence's material appears to be related to the ideas about how the syntactic structure of a sentence suggests its articulation into given and new. Diesing's proposal, however, applies at LF, that is, at a level of representation at which the scope of operations has already been assigned, thus cannot be interpreted as a proposal for assigning a scope to operators

## 6.3 ARGUMENT SELECTION, THE THEMATIC HIERARCHY, AND THE EVENT DESCRIPTION BUILDING PROCEDURE

In §1.5 I proposed that the preference conflicts in passive sentences observed by Karttunen and MacDonald, as well as Kuop's data about the wide scope assigned to NPs filling certain grammatical functions, are the result of the interaction between the given/new contrast and a second interpretation procedure, the process that maps grammatical relations such as subject and object into the arguments of a predicate, variably called *linking theory* or *argument selection* [Dowty, 1991]. I have discussed my representation for roles and predicates in §5.5. I now discuss argument selection and its effect on scope in more detail

### 6.3.1 Argument Selection and the Thematic Hierarchy

It has been proposed<sup>12</sup> that the facts about argument selection illustrated by (5 26) can be accounted for by stipulating that thematic roles are ordered according to a *thematic hierarchy*. This order usually looks like the form

AGENT < EXPERIENCER < < LOCATION < SOURCE < GOAL < THEME

The effect of the thematic hierarchy on argument selection is specified by the following principle

<sup>11</sup> The original idea may be due to Karttunen [Karttunen, 1976]

*Argument Selection Principle:* In predicates with grammatical subject and object, the argument that fills the role higher in the Thematic Hierarchy will be lexicalized as the subject, the argument that fills the next higher role in the TH will be lexicalized as the object.

What's interesting about the Thematic Hierarchy idea (as opposed to Kuop's proposal about a Grammatical Function Hierarchy) is that there seems to be independent motivation for it. Two phenomena that provide evidence for such a hierarchy are discussed by Jackendoff in [1972]. Jackendoff observes that not all actives have passive counterparts, and the disparity is related to the thematic roles of the predicate to be passivized. (6 15a), for example, is ambiguous between two readings. In one of these the subject is an agent; this reading can be paraphrased as "Tom used a brush on the wall". In the other reading, the subject is a theme: "Tom's body inadvertently brushed against the wall". When the expression is put in the passive, however, as in (6 15b), only the first reading is maintained.

- (6 15) a Tom brushed the wall.  
b The wall was brushed by Tom.

Verbs of measurement do not passivize at all, as shown by the examples in (6 16) and (6 17). In these sentences, the D-structure subject is the THEME, and the measure phrase indicates a LOCATION on a scale; the sentence specifies the position of the THEME on the scale provided by the LOCATION.

- (6 16) a The book costs five dollars.  
b \*Five dollars are cost by the book.  
(6 17) a Bill weighs two hundred pounds.  
b \*Two hundred pounds are weighed by Bill

Jackendoff considers next the contrast between the pair of sentences in (6 18) and the pair in (6 19). In (6 18a), "John" is the GOAL, "the letter" is the THEME. This may be passivized, as in (6 18b). However, when the D-structure subject is the THEME, and the D-structure object the GOAL, as in (6 19), the sentence cannot be passivized.

- (6 18) a John received the letter  
b The letter was received by John.  
(6 19) a John reached the corner.  
b ??The corner was reached by John

These data are predicted by the Argument Selection Principle. The b cases in (6 16), (6 17) and (6 19), as well as the unpassivable reading of (6 15b), are all sentences in which the NP in Subject position lexicalizes an argument that fills a role that is lower in the TH than the role filled by the argument lexicalized by the NP as object position.

According to Jackendoff, the thematic hierarchy plays a role in reflexivization, as well. He notes the contrast in (6 20). Why is (6 20b) much worse than (6 20a)? (Note that neither violates Principle A of the Binding Theory.)

- (6 20) a I talked to John about himself  
b ??I talked about John to himself.

Reflexives aren't allowed as the VP in the *by*-phrase of passives, either, as shown in (6.21):

(6.21) \*John was shaved by himself

Other cases of bad reflexivization are ruled out by Principle A were discussed by Postal (1971); the sentences in (6.22) are examples of the constraint in reflexivization in cases of *themselves* movement.

- (6.22) a It is tough for Tony to shave himself  
b \*Tony is tough for himself to shave  
c \*Himself is tough for Tony to shave

The generalization proposed by Jackendoff in account for these facts also makes use of the Thematic Hierarchy.

**Thematic Hierarchy: Condition on Reflexives.** A reflexive may not be higher on the Thematic Hierarchy than its antecedent (Jackendoff 1972, p. 148).

Dowty proposes the following principle to account for the  $\pi \in$  of Proto-roles in argument selection.

**Argument Selection Principle (Dowty 91).** In predicates with grammatical subject and object, the argument for which the predicate entails the greatest number of proto-agent properties will be lexicalized as the subject of the verb; the arguments having the greatest number of proto-patient entailments will be lexicalized as the direct object.

Two corollaries of this principle are, first, that if two arguments of a relation have approximately the same number of entailed proto-agent and proto-patient properties, either or both may be lexicalized as the subject. Second, with a three-place predicate, the non-subject argument having the greater number of entailed proto-patient properties will be lexicalized as the direct object, while any other non-subject argument will be lexicalized as an oblique or prepositional object. The principle does not require proto-roles to classify all arguments (some arguments have neither role); it allows for two arguments to share the same role, and some arguments may qualify partially, but equally for both roles.

If we assume the Argument Selection Principle and if we take proto-agent and proto-patient to be, 'in the extreme', on a line from 'most likely to occur as subject' to 'least likely to occur as subject', we obtain a partial order among roles. In other words, the thematic hierarchy follows naturally from Dowty's account of semantic roles without the need of further stipulations.<sup>12</sup>

<sup>12</sup> In truth, the hierarchy one obtains for an Event's functions depends on what one takes to be the order relation, and in fact the hierarchy proposed by Dowty is slightly different. There are several versions of Thematic Hierarchy, not all of which are used by Jackendoff. The simpler version is the one used by Dowty and King (p. 1) that only requires a simple role (Agent) to be distinguished.

### 6.3.2 An Event Description Building Procedure

When building the model of the situation described by a sentence, a listener may try to simultaneously determine the fillers of all the roles, or may try to identify one role after the other: I propose that the latter is the case. I propose, that is, that the order between roles specified by the Argument Sels. Non Principle provides listeners with a default procedure for building event descriptions, in the sense that a listener starts builds first those parts of the event description that are related to the fillers of role types higher in the hierarchy. I call this the Event Description Building Hypothesis.


**Event Description Building Hypothesis (EDBH):** The thematic hierarchy specifies the default order in which an event description is built.

The EDBH serves as the basis for a default procedure for building a description of the situation described by an utterance. This procedure consists of weak default rules that implement argument selection. These rules essentially perform generalized application as at the underspecified representation level; the EDBH is implemented by assuming that these rules must apply in a certain order. In the case of prepositional quantifiers, for example, I assume that there is a model construction rule schema like MCR PRE-SUBPP AGENT below, that is triggered by the occurrence of the function

$$\lambda P [\text{def } z (x \models Qz) // P]x$$


(where *def* is a prepositional determiner) in the position of [Spec TP] when the main verb is a reflexive that requires an AGENT role, and that produces a new hypothesis by 'applying' the interpretation of the NP to the rest of the underspecified representation—which, in practice, means applying the semantic translation of the NP to the property obtained by replacing the lexical interpretation of the NP with a variable.

<sup>13</sup> This operator is defined in §5.5.



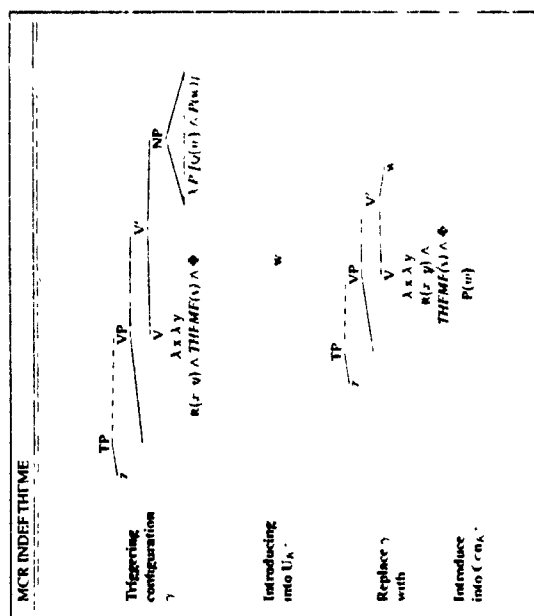
The result of the rule for the non-presuppositional reading of indefinite NPs, instead, is to add new discourse markers to the current DRS instead of introducing a tripartite structure, as follows:<sup>14</sup>

<sup>14</sup> This rule works for indefinite NPs such as "a horse" or "some horses"; a similar rule for cardinal NPs such as "three horses" – and perhaps also for the non-prepositional reading of NPs such as "many horses" – can be achieved by treating the determiner as a predicate stating the cardinality of a set and adding to the constraint grammar an additional condition of the sort  $\text{deter}(z)$  where  $z$  is the discourse marker introduced by the cardinal NP.



The rules above are very similar to the model construction rules proposed by Kamp and Reyle (1993) (compare, for example, *MCR-PRESUPP-AGENT* with Kamp and Reyle's *CREVERY*, at pag. 169 of (Kamp and Reyle, 1993)); in fact, one could claim that the rules proposed above are already based on an implicit ordering among arguments. There are important differences, however: first of all, *MCR-PRESUPP-AGENT* and *MCR-INDEF-AGENT* are weak defaults, and therefore they apply only if no rule with a higher priority has resulted in a 'discharge' of the NP. Secondly, they depend for their application on a predicate's looking for an agent, whereas Kamp and Reyle's rules are independent from the lexical semantics of the predicate.

The ordering between thematic roles is implemented by forcing the argument selection rules for roles 'lower' in the thematic hierarchy to 'wait' until the MCR rules building the part of the event description related to 'higher' roles have been executed. This is done by requiring the trigger of rules such as *MCINDEX:THEME* (the equivalent in this system of the rule for indefinites in object position in [Kamp and Royle, 1993]) to have an object of type *e* in subject position, instead of a property of properties.



### 6.3.3 Scope and the Thematic Hierarchy

The EDBH provides an explanation for several scopings preferences that have been observed in the literature, yet cannot be accounted for if only the effect of the given/new partition discussed in §6.2 is assumed.

First of all, Joupp motivated her Grammatical Function hierarchy by observing that, in addition to a preference for the Subject to take wide scope over the Object NP, there is a similar preference for the Indirect Object to take wide scope over the Direct Object, and for the Prepositional Object to take wide scope over the Direct Object. The first preference is illustrated by the examples in (6.23). The preferred interpretation of (6.23a) is the one in which "every child" takes wide scope; this preference is not changed by reversing the order of the quantifiers, as in (6.23b), nor by switching the quantifiers, as in (6.23c) and (6.23d), whose preferred reading is again the one in which the NP in Indirect Object position, "a child", takes wide scope. These d's are all in agreement with the predictions of the EDBH as well, since the Direct Object of the predicate *TELL*, a THEME, is lower in the hierarchy than the Indirect Object (a GOAL).

- (6.23) a I told every child a story.  
b I told a story to every child.

- c I told a child every story.  
d I told every story to a child.

The preference for the NP in Prepositional Object position to take wide scope with respect to the NP in Direct Object position is illustrated by the sentences in (6.24).

- (6.24) a I had many conversations with a friend.  
b I had a conversation with many friends.  
c Freddy hit many balls with a bat.  
d Freddy hit a ball with many bats.

Another phenomenon which cannot be explained in terms of the Given/New contrast only is the lack of a clear preference for passive sentences—i.e., the contrast between the clear preference of the subject for assigning wide scope to "a kid" in (6.25a), while at the same time finding (6.25b) ambiguous.

- (6.25) a A kid climbed every tree.  
b A tree was climbed by every kid.

This contrast cannot be explained by stipulating a Grammatical Function hierarchy, either, since that would predict no contrast between (6.25a) and (6.25b). What seems to be happening is that two contrasting principles are at work: one assigning a wide scope to the subject NP, the other assigning wide scope to the object NP. By assuming the EDBH in addition to the Given/New contrast, we explain the conflict.

Finally, Kurtzman and MacDonald's result that their subjects have a much stronger preference for the NP in subject position to take wide scope when the NP is weak may be explained, as follows. Once MCR NON-PRESUPP AGENT is applied, a participant to a conversation has nothing more to do. However, after MCR PRESUPP AGENT is applied, the c.p. still has to identify the resource situation of the quantified NP; so, at least one more "weak" inference step is required. The interpretation thus obtained is weaker than the interpretation obtained after applying MCR NON-PRESUPP AGENT.

### 6.3.4 Additional Predictions

If we assume Dowty's treatment of thematic roles, and if we assume the EDBH, we get additional predictions that would be interesting to verify. Notice that I have formulated the rules above in terms of AGENT and THEME roles. I assume, that is, that the application of a rule such as MCR PRESUPP AGENT involves a step during which the listener evaluates the entailments of the predicate occurring in the sentence; the rule is applied only if the listener thinks that there is a good enough match. Thus, the more 'agentive' a predicate is, the easier it is to apply the rule. All other things being equal, we predict that since basic transitive verbs have prototypical AGENTS and prototypical THEMES, the preference for the subject NP to take wide scope over the object NP should be stronger than for symmetric predicates such as "meet", in which the two roles do not differ much in prototypicality, thus the only preference is due to topicality effects.<sup>15</sup>

<sup>15</sup> This idea actually assumes that different effects may sometimes combine.



The data are not clear, but there seems to be a slight contrast between the (a) and the (b) sentence in (6.26) (although this contrast could be due to world knowledge as well)

- (6.26) a Every student built a model airplane  
b Every student met a professor

For the same reason, there should also be a contrast in scoping preferences between basic transitive and psychological predicates

- (6.27) a I very student built a model airplane  
b Every student feared a professor

Thematic hierarchy effects should be weaker for psychological predicates that 'come in pairs,' such as to fear/to frighten. This is in agreement with my intuitions about (6.28)

- (6.28) a I very student feared a professor  
b A professor frightened every student.

Finally, in commercial transaction predicates, the seller fills simultaneously both the AGENT and the SUBJECT role, while the buyer fills simultaneously both the AGENT and the GOAL role. One may argue that this, again, should weaken the thematic hierarchy effects, but whether this is actually the case is not clear to me

- (6.29) a A student bought a model airplane from every ship  
b Every student bought a model airplane from a ship  
c A student sold a model airplane to every ship  
d Every student sold a model airplane to a ship

## 6.4 THE INTERPRETATION OF OPERATORS

In the rest of the chapter I am going to discuss the most important among the discourse interpretation procedures that play a role in the TRAINS conversation. The interpretation of all operators involves two kinds of rules. First of all, there are defeasible principles formulating hypotheses about the interpretation of contextually dependent aspects of the meaning of the operators. I call these principles for anchoring resource situations, for reasons that should become clear below. Second, there are model construction rules, that are like traditional DRT construction rules except that they need to be 'enabled' by a principle for anchoring resource situations

For each operator I first present the model construction rule(s)—i.e., I specify how that operator contributes to the common ground, and which conditions have to be met in order for that contribution to occur. Next, I discuss the discourse interpretation procedures that lead in these conditions being satisfied

## 6.5 DEFINITE DESCRIPTION INTERPRETATION

The interpretation of definite descriptions is one of the most important processes that occur in our dialogues, and the one whose interaction with other interpretive processes such as the cum-

positions of the given/new partition, or the default procedure for building situation descriptions, is most clear

As discussed in §3.7.1, the two most common cases of definite descriptions in the TRAINS conversations are anaphoric definites and definites interpreted with respect to the visual situation. In both cases, interpretation involves both 'focusing' inferences and principles for choosing an interpretation. (I have discussed the location theory of definite descriptions, on which the following proposal about the interpretation of definites is based, in §5.3.) The principle accounting for the interpretation of visible situation use of definite descriptions states that when both participants to a conversation are focused on the same part of the visual situation, that part of the visual situation may be used as the resource situation for definite descriptions. The principle accounting for the interpretation of anaphoric uses of definite descriptions states that the resource situation of definite descriptions may be identified via the current discourse topic.

In the rest of this section, I discuss first the model construction rule for definite descriptions, then the two interpretation procedures, and finally the effects of definite description interpretation on scope

### 6.5.1 Model Construction Rule

The model construction rule for definite descriptions, MCR DEF, is presented below. This rule schema is enabled by prior identification of the resource situation  $s$  of the definite, as specified by the location theory (see §5.3). As in standard DRT, definites do not introduce complex conditions in the current situation description, unlike standard DRS, however, the situation description containing the logical form in which the NP is included is modified, rather than the root DRS. A new situation description is however added to the content of the current situation descriptions, specifying that the resource situation  $s$  of the definite description includes  $y$  as the unique object of type  $P$ .

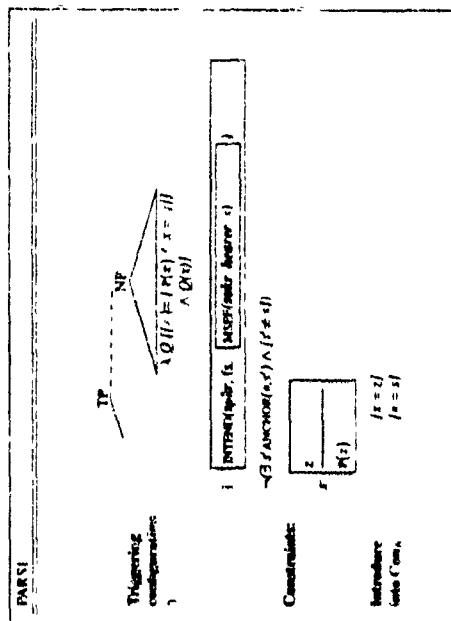
(191)

29 1 U okay.  
29 2 great  
29 3 while this is happening.  
29 4 take engine E1 to Danville  
29 5 pick up the benzene  
29 6 and come back to Area  
30 1 S okay

As mentioned above, two bands of inferences seem required to interpret this kind of definite descriptions. First of all, a conversational participant must know that if the visible stimulus contains an object of the appropriate type, that object can be the intended referent for the definite description. The principle for anchoring relevance situations presented below, PARSI, says just that:

**PARTS1** If a speaker uses a referring expression "the *x*", the speaker intends the mutual attention of the conversational participants to be focused on the situation *s*, and the visible situation contains an object of type *x*, then the listener may hypothesize that *s* is the resource situation for "the *x*".

in the following table. The results are summarized in Table 1.




In the formulation of PARS3 above, I have assumed that 'parameter anchoring' is done by adding statements of the form  $[p = sf]$  to the DRS in which the logical form occurs, and I have ignored temporal issues, such as the fact that the intention concerning the mutual focus of attention should be at the same time as which the utterance of the definite description occurs. This latter problem could be easily fixed by complicating slightly the representation, as it is indeed the case in the actual implementation.

As discussed in §3.7.1, the visible situation use of definite descriptions depends on the current focus of attention when an object is in the current focus of attention, it can be felicitously referred to by means of a definite description even when other objects of the same type have been introduced in the discourse or are part of the world described by the map. An agent's visual focus of attention changes continuously (Allport, 1987), yet not all of these shifts can be exploited to make the use of a definite reference felicitous, since conditions on mutual knowledge have to be met (Clark and Marshall, 1981). An attention shift can only be exploited when the participants in a conversation mutually know that the shift took place, on the grounds of some general fact about the conversation. In the case of the 'air commander' conversations studied by Ginzburg (1977), the movement of the focus of attention was related to the structure of the task. In our transcripts, the most important conversational principle governing visual attention shifts appears to be the following:

**Follow The Movement.** Part of the intended effect of an utterance instructing an agent to move an object from one location to another is to make the terminal location of the movement the new mutual situation of attention.

This principle can be formalized by the following (strong) default


<sup>19</sup> This principle could perhaps be derived from a more general principle saying that in our diagrams it is expected of each conversational participant that he pays attention to what the other participants say, and paying attention when the movement of an object in the map is connected simply ensures following the movement of that object on the map. This however is not very important for our purposes.



The system can be read as follows: the occurrence of a conversational event  $ce$  of type  $INSTRUCT$ , where the restriction is one to move the object  $z$  to  $p$ , results in hypothesizing that the discourse situation includes an intention  $i$  of the speaker  $s$  that the mutual focus of attention of the speaker and the hearer be the situation  $PLACE(p, M, pS)$ , which is the 'place' situation type consisting of the information in the map about position  $p$ . This state  $fs$  of the mutual attention being on  $PLACE(p, MapS)$  holds until the next conversational event in discourse segment  $ce$ .

Let us see now these principles at work in the interpretation of utterances 29.4-29.5 in (1.61). The conversational event generation rules assign to 29.4 the translation in (6.30).


29.4 U take engine E1 to Denzville.



<sup>20</sup> This is a simplification: actually, one should 'come out' the place situation not from  $MapS$ , but from the situation 'working'. We're aware of moving  $z$  to  $p$ . Also, it's not clear whether a more general formulation in terms of attention rather than... surface act would hold.


In (6.30),  $c_1$  is the conversational thread of which the new conversational event,  $ce_{29.4}$ , is part. The conversational event takes place at time  $P$ , and is an instruction of moving engine E1 to Dansville.

Assuming a simple inference to the effect that each 'taking' event in the TRAINS domain entails a move, the 'follow the movement' principle expressed by rule FTM can now apply, with the result that the listener hypothesizes that the user intends the 'place situation' consisting of the facts about Dansville to be the  $loc$  mutual situation of attention. This hypothesis is represented as follows:




The next utterance, 29.5, is interpreted in the common ground thus augmented. The conversational event associated with 29.5 is described in (6.31).

29.5 U pick up the boxcar




Now the principle for anchoring resource situations mentioned before, PARS1, may apply, with the result that the 'system' has hypothesized that the 'user' intends the resource situation for the definite boxcar in 2.3 to be  $loc$  (Dansville MapS).



As there is only one boxcar in  $PLA$   $loc$  (Dansville MapS),  $h$  it is consistent to infer (6.32). This entails (6.34).<sup>71</sup>

<sup>71</sup> It appears that persistence reasoning may also be part of a listener's interpretive rules, i.e. a novel user input in the form of an instruction that the previous focus of attention should be maintained. For example, such a user input may be used as



### 6.5.3 Anaphoric Uses of Definite Descriptions

Principles for anchoring resource situations and focusing principles also constitute the basis of the processes involved in the interpretation of definite descriptions used anaphorically.

In the case of anaphoric definite descriptions, the focusing element is the discourse topic, as discussed in §3.7 and §3.8. A proper discussion of discourse topic change requires, however, a much more extensive discussion of conversational events and the intentions they convey that I can include here. I will simply say that both principles enforcing the effect of discourse signals (Reichman, 1985; Schiffrin, 1987) and 'persistence principles' are needed, the latter being weak construction rules that suggest that unless some stronger defaults suggest otherwise the conversational thread doesn't change.

The principle for anchoring resource situations involved in the interpretation of anaphoric uses of definite descriptions such as 'the boxcar' in 5.1 (from the transcript in Fig. 1.2, §1.4, repeated below)

3.1 U now when so we need to get a boxcar in Corning.  
3.2 where there are oranges  
3.3 there are oranges at Corning  
(6.35) 3.3 right?  
4.1 S right  
5.1 U so we need an engine to move the boxcar  
5.2 right?  
6.1 S right

is called PARS2


PARS2 If the current discourse topic is the situation or situation kind  $s$  that includes a discourse marker  $z$  of type  $P$ , a definite NP of the form 'the  $P$ ' may be taken to refer to  $z$  if the NP is lexically primed by  $z$ .

This principle can be formulated as a construction rule as follows (I have simplified the representation of the trigger for the rule, only including the recurrence of a NP of the appropriate type inside a TP).

a full back case as examples such

(6.31) Send the engine to Elmore, and send the boxcar with it.

It seems that this inference can be used in case the current focus of attention is on a boxcar on the map at which there is a boxcar, the first part of the utterance is a moving selection, and therefore the FTM would suggest a shift of the focus of attention to Elmore, however, this interpretation is overridden when interpreting the second part of the utterance, as the predicate 'send' with 'with' requires the boxcar to be at the original location of the engine. Several complete inferences need to be resolved before incorporating this form of focus in the present system including making sure that the focus of attention is not always automatically shifted.



As discussed in §1.5,  $\alpha$  make the assumption that lexical priming relations are established and that this information is available during discourse interpretation in the form of predicates of the form  $LPRIME(P, Q)$ , both psycholinguistic and computational models of priming exist and any of them will do the job required here

#### 6.5.4 Definite Description Interpretation and Scope Disambiguation

The effect of definite descriptions if interpretation on scope disambiguation is demonstrated by the examples in (6.35), in which the weak defaults presented in §6.3 and §6.2 are overridden. The preferred reading of (6.36a) is the one in which a single orange factory gets reached by many boxcars, this is not the reading that would be obtained by the EDBP, that would assign to the NP filling the agent role wide scope over the NP filling the theme role. Similarly, in (6.36b), the NP "the engine," the THEME of "hook," takes wide scope over the NP "each boxcar," the GOAL of the verb, again overriding the default preference provided by the thematic hierarchy

- (6.36) a Many boxcars reached the orange factory  
b Hook the engine to each boxcar

These preferred readings are, according to the theory presented here, the results of the fact that whether the definites in (6.36) are anaphoric or interpreted with respect to the visual situation, the processes that assign an interpretation to definite descriptions take precedence over the weak default provided by the EDBP. The preferred extension of the theory is thus obtained


by first applying the actions formalizing the discourse interpretation procedures associated with definite description interpretation, then the weak defaults (see §3.9). When these latter apply, the interpretation for the definites has already been determined. This assumption about the priority of definite description interpretation processes over other interpretation processes is not a stipulation introduced for the sake of explaining the facts about scope, but it's a generally accepted fact about interpretation (see, e.g., Crain and Steedman's work on the effect of reference on parsing preferences [Crain and Steedman, 1985]).

The DRT construction algorithm proposed in earlier drafts of [Kamp and Reyle, 1993] produces the same results than the algorithm discussed here for (6.36), but different results when narrow scope definites are considered.<sup>22</sup> This can be shown by looking at examples like (6.37), in which the interaction of discourse interpretation with scope disambiguation plays a crucial role.<sup>23</sup> The older version of the DRT construction algorithm generates for this sentence an interpretation in which there are a single principal and a single meeting, as opposed to the preferred reading in which there is a single meeting, but a different principal per school.<sup>24</sup>

- (6.37) Every school sent the principal to the meeting.

According to the model of interpretation of definites just discussed, the interpretation of the sentence proceeds as follows. Assume that the context  $\gamma$  (6.37) is sentence (6.38), that introduces into the model the situation description in (6.39), containing a group of schools and a meeting.

- (6.38) There was a meeting of the schools in the district.



(6.39)


Since  $x$  is the discourse topic, PARS2 can be used to identify  $\alpha$  as the resource situation of the NP "the meeting". It seems plausible to assume that a principle similar to PARS2 exist, capable to establish an anaphoric connection between the NP "every school" and the discourse marker  $S$ . No unique principal can be identified in  $x$ , however. The model construction rules MCR DEF for definite descriptions and MCR EVERY for universal quantifiers (discussed in §6.8, but not that much different from the rule proposed by Kamp and Reyle) can apply after

<sup>22</sup>As mentioned in §1.5, the version of the rule found in the final version of the book does yield the preferred interpretation for narrow scope definites. On the other hand, it does not produce the preferred interpretation for (6.36).

<sup>23</sup>This example is a variant of an example from [Harnad and Scheibel, 1986].

<sup>24</sup>While  $\alpha$  may seem strange, narrow scope definites are not really that rare, especially when adverbs of quantity are considered. The literature includes examples such as "If a dog barks at a cat, the cat always meows" or "When a judge finds a person guilty the person has ten days to appeal the sentence." find.

discourse interpretation, which results in the partially disambiguated interpretation in (6.40), in which the resource situations of "every school" and "the meeting",  $s_1$  and  $s_2$ , have been identified, and where the "extended" notation for universal quantifiers introduced in §4.3 is used (I also represented the definite NP "the principal" among quotes to save space).



(6.40)


Note that in (6.40) a situation including a school and a marker for it have been introduced ( $s_1$  and  $s_2$ , respectively). Now the rules for interpreting associative definite descriptions may apply which I assume are going to be like PARS2 in that they are activated by lexical priming relations between a discourse marker and an as yet unresolved definite NP. Such rules can establish relations between a relational predicate such as PRINCIPAL and a priming predicate such as SCHOOL. The result is that  $z$  is chosen as antecedent of  $x$ , and  $s_2$  is chosen as the resource situation for the definite "the principal". Once this is done, MCR DEF may apply and update the description of  $s_1$ , accordingly.<sup>25</sup>

An interesting case of interaction between definite description interpretation and scope is (6.41). The preferred reading of this sentence is the one in which a single tree gets climbed by every kid. To get this reading, the default preference for NPs filling the agent role to take wide scope over NPs filling the theme role must be overridden as in (6.36), however, it doesn't seem necessary for the context to include a tree in order to get this reading—in other words, it doesn't seem necessary for there to be a unique tree in the common ground to which "the tree" refers in order to get the preferred interpretation. If there isn't such a tree, the sentence is perceived as pragmatically infelicitous, but the preferred reading is still the one in which "the tree" takes wide scope. Why is this the case?

(6.41) Every kid climbed the tree

<sup>25</sup> From what I just said about the interpretation of (6.37), one may infer that in order to get the facts right it is necessary to assume "at the interpretation of 'the principal' is delayed until an appropriate situation is added to the common ground. While this may actually be the case, it would be preferable (although messy) to develop a version of PARS2 and of the rule for associative interpretation that deal with cases like the one just presented in a principled way by working directly on the logical form.


My answer is that the constraints on the resource situation of definite descriptions discussed above play a role here. If there is no tree in the common ground which can be interpreted as the referent of the definite descriptions, the discourse interpretation procedures described in this section do not apply. So, an interpretation can only be obtained by means of the given/new procedure and the EDIP. But after rules such as SUBJECT GIVEN or MCR PRESUPPOSITION AGENT have applied, we still obtain an interpretation which contains a presupposition: NP, i.e., the familiarity condition still has to apply.



(6.42)

A process of accommodation has to take place—a new situation is hypothesized to be part of the common ground, at the top level, and the situational parameter representing the resource situation of the definite is anchored in the new situation. Cases like (6.41) thus seem to require a third level of defaults, including accommodation procedures of this sort.

The reading represented by (6.43), instead, is not available. (I have ignored situation descriptions, etc.) Note that this constraint is automatically satisfied by the structures obtained by the procedures described in this section.



(6.43)

The interaction between definite description interpretation and the given/new partition is displayed in (6.44), where two discourse interpretation procedures with the same priority apply one that makes the existence of a set of towns part of the presuppositions of the sentence, and one that assigns an interpretation to definite descriptions.

(6.44) In every town, the newspaper interviewed Kissinger

If a newspaper is part of the common ground, the two procedures apply at the same level and their results may be merged because they are not inconsistent; the sentence is thus perceived as not ambiguous. This is the case in (6.45), for example.

- (6.45) The Wall Street Journal pays a lot of attention to the opinions of Republican politicians.  
In every town, the newspaper interviewed Kinsing-7


If, however, no prior context is given, only the given/new procedure applies at first. This interpretation can be extended in complete interpretation in two ways by using accommodation procedures: either by establishing a weak associative relation between "the newspaper" and "every town," or by introducing a single situation in the continuous ground, that includes a newspaper that "the newspaper" can be interpreted as anaphoric on. My informants agree that the sentence in this case is perceived as ambiguous.

## 6.6 THE INTERPRETATION OF MODAL AUXILIARIES AND MODAL VERBS

The interpretation process for modals, like the interpretation processes for definite descriptions and for tense verbs discussed, can be characterized by two distinct sets of rules: principles for determining the modal base that resolve the contextual component of meaning—in the case of modals, the modal base, and modal construction rules that rewrite a partial interpretation when the modal base has been identified. I will discuss these two kinds of rules, in turn.


### 6.6.1 Modal Construction Rules

The modal construction rule for modal auxiliaries is an axiom schema called MCR MODAL-AUX. In the schema,  $M$  is a modal, and  $X$  can be either VP or NEGVP. The rule is triggered by the occurrence of a logical form in which the left position is occupied by a modal auxiliary. The rule is enabled whenever previous discourse interpretation processes have identified the modal base of the modal, that is, have hypothesized in advance for the parameter ranging over situations kinds that is part of the lexical meaning of the modal. The schema results in the replacement of the logical form in the trigger with a complex condition that contains in the nuclear scope a logical form obtained from the trigger by eliminating the modal from the head position.



As discussed in §5.9, I assume that in the case of modal auxiliaries, the interpretation of tense (specified as tense' in the rule) is merged with the interpretation of the modal.

The modal construction rule for modal verbs is called MCR MODAL-VERB. It is triggered by the occurrence of an utterance whose logical form includes a modal verb as head of a VP, and a tenseless TP as complement of that same TP. The result of MCR MODAL-VERB is similar to that of MCR MODAL-AUX, except for the difference in the syntactic tree that gets rewritten.



### 6.6.2 Discourse Interpretation Rules for Models

In the conversations in the TRAINS corpus, sentences containing modal auxiliaries and modal verbs are used, first of all, to specify the task the user and the system have to accomplish at the beginning of a conversation. In this use, modals like "need", "have to", and "better" are semantically interpreted as expressing deontic obligations; pragmatically, they are interpreted either as providing information to the other party, point, or as suggestions. Examples include:

- (6.47) a U: We have to make OJ.  
 b U: We need to get a boxcar of oranges to Bath by 8 AM.  
 c U: I need to get oranges to Bath by 8 AM.  
 d U: We better ship a boxcar of oranges to Bath by 8 AM.  
 e U: I have to get one tanker of OJ to Avon and a boxcar of bananas to Corning, uh by 3 PM.

Modally qualified sentences, are also used in the second phase of typical TRAINS dialogues, during which the user and the system discuss the constraints of the TRAINS world and consider subgoals that have to be achieved in order to accomplish the task. The user's utterances in this context play the pragmatic role of "checks", in the sense that the user uses the system to verify facts within possible plans. Examples include:

- (6.48) a U: So we need an engine to move the boxcar.  
 b U: Two engines can't run on the same track, can they?

<sup>10</sup> There are only two exceptions to the rule that the task is specified by means of a sentence that contains a modal verb:  
 a U: one tanker car of orange juice to Avon and a boxcar of bananas to Corning, both by 5 PM  
 b U: we've got a more complicated problem.

Finally, modals like "should" are used by the user to formulate proposals concerning the plan and/or by the system to formulate suggestions (the system's suggestions are at one constraint only expressed using the modal auxiliary "could"). These utterances have the pragmatic force of proposals.

- (6.49) a U: We should hook up an engine to the boxcar.  
 b S: We could use one engine to move both the boxcar and the tanker car.

The process involved in the interpretation of modals can be illustrated by discussing the interpretation of the following sequence of utterances.


- (6.50) a U: We should use a boxcar to move the oranges.  
 b U: We should hook up an engine to the boxcar.

The basic step of modal interpretation is the identification of the modal base. This process is performed by default inference rules, much like those that anchor the resource situation of deontic descriptions. For example, the rule PAMB SHOULD presented below generates an hypothesis according to which the modal base of the modal SHOULD is determined by the situation kind associated with the discourse topic of the current segment in the case of TRAINS conversations, the topic situation usually is the "proposed plan" (Plan<sup>9</sup>). The rule may be paraphrased as follows:

PAMB SHOULD: Take the modal base of a modal statement to be the situation kind associated with the set of propositions representing (what is known about) the current discourse topic, if it is consistent to do so.

This discourse interpretation rule is specified as follows. The rule is triggered by the occurrence as part of the description of situation  $s_1$  of an underspecified logical form representing a TP whose head is the translation of the modal "should". If the parameter  $k$  in the translation (representing the situation kind that determines the modal base) has not yet been anchored, and the topic of the current discourse segment  $t_1$  is the situation or situation kind  $s$ , then produce an hypothesis in the effect that the description of situation  $s_2$  is augmented by having the value of the parameter  $k$  to (the situation kind associated with)  $s$ .






Consider for example (6.50a), and let the current topic of conversation be the situation kind "plant". PAMB SHOULD results in the hypothesis that the modal base of the modal "should" in (6.50a) is 'what is known about the plant', that is, (6.50a) gets interpreted as "In view of what is known about the plant, we should use a vasciar to move the oranges" 37.

Once PAMB SHOULD has applied, MCR MODAL-AUX can apply as well, resulting in the updated hypothesis in (6.51). The event description that triggered PAMB SHOULD is replaced in the characterization of  $s_2$  by a new event description in the form of a complex condition with the form *should*(K<sub>1</sub>)(K<sub>2</sub>).

<sup>27</sup> In our dialogues, *it* which the context is relatively fixed, simple rules such as PAMB SHOULD suffice to determine the model here. In order to handle Kravetz's examples of modal base choice more complex procedures of fact communication would be required so which the situation determining the modal base has to be reconstructed on the basis of the fully interpretation of the sentence.



(159)

### 6.6.3 Plans de Dissuade Tendres

The fundamental contribution of dynamic theories of discourse such as DRT has been to assign to sentences an interpretation which can be used to incrementally construct the characterization of a situation/discourse topic out of the interpretations assigned to the sentences in a text. I discussed in Chapter 4 how situation descriptions are used in CRT to preserve this aspect of DRT while at the same time allowing for a more indirect relation between the common ground and the situation described by a text or conversation.

What about plans, however? The TRAINS conversations show that a plan can also more-or-less be modally qualified and imperative intertext, that is, that the plan can also serve as a discourse topic. A plan can also serve as a discourse topic in the sense that objects introduced in the common ground when describing aspects of the plan by means of imperative intertext can be felicitously referred to when discussing the plan.

Let's first the question of how conversational participants get from (6.51) to an update of (6.52). Participants such as 'We should use a boxcar to move, the oranges' only interfere with the common ground portions of a plan.<sup>24</sup> My suggestion is that once (6.51) is obtained, the hearer makes the further inferential step of applying a defeasible rule for speech interpretation. These rules determine the preferred interpretation of the utterance in context, as in the spirit of Perraelli's proposal (Perraelli, 1990), the relevant rule may be paraphrased as follows:


**Model to Proposal:** A model statement of the form "in view of the following facts about S, we should do X" can be taken to express a proposal to add an action of type X to the planned course of action, if it's consistent to do so.

am suggesting, in other words, that when entering (6.50a), the user is proposing to make the disacourse situation one in which *Plan*<sup>9</sup> is a subkind of the kind of situation in which oranges are moved by means of a balancer. As we shall see, this hypothesis not only accounts for the

<sup>20</sup> An issue I won't touch here is the fact that the plan is not 'put together' immediately; the participants go back checking whether the plan is coherent and achieves the intended goal only after having introduced several of its requirements, especially when the conversation is long.

fact that plans can be constructed incrementally, but also provides a way to deal with modal subordination

The hypothesis about (6.50a) obtained by this rule is as follows:



The interpretation of the modal in (6.50b) proceeds as the interpretation of the modal in (6.50a)

#### 6.6.4 Modal Subordination

The antecedent of the definite "the buscar" in (6.50b) is the indefinite "a buscar" introduced by (6.50a). This is an example of modal subordination (Karttunen, 1976; Roberts, 1987). This phenomenon is illustrated by the contrast in (6.54) while anaphoric reference to discourse markers in the scope of modal operators is in general ruled out, it becomes possible if the anaphoric expression is in the scope of a modal operator as well.<sup>20</sup>

- (6.54)
- a We need to find an engine. \*If/The engine is powerful
  - b We need to find an engine. If/The engine should be powerful.
  - c There is an engine at Avon. We should rent it/the engine at Bath

The approach to modal subordination pursued here derives from ideas in [Jackendoff, 1972], as elaborated by Roberts [1987] and in [Pesveto and Zucchi, 1992]. Jackendoff observes what he calls *conference condition on modal dependence* ([Jackendoff, 1972], p. 294)

*Conference Condition on Modal Dependence:* If NP<sub>1</sub> and NP<sub>2</sub> are intended to be *conferential*, they must be dependent on

<sup>20</sup> In the TRAINS conversations we also find examples like the following, showing that constructions other than modals (e.g. imperatives) can be used to establish the form of subordination that is necessary to make anaphoric reference grammatical

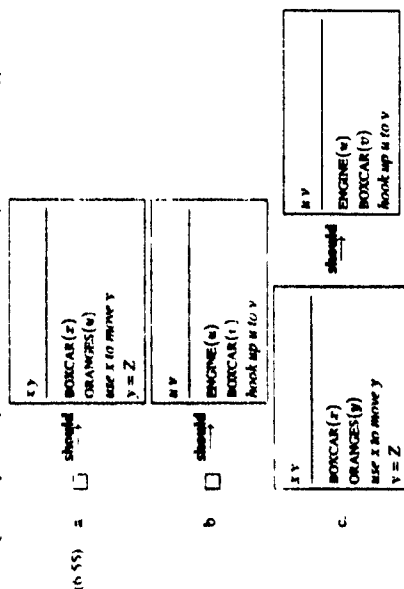
- (6.51)
- a Let's send an engine to Avon. \*It's broken to the buscar
  - b Let's send an engine to Avon. We should bank it to the buscar

Although I will not discuss modal subordination with imperatives here, the proposal I make can be extended to deal with the latter as well

- a. the same type modal operators (weak form)
- b. the same robust modal operators (strong form)

According to Jackendoff, the intuitive motivation for the *Conference Condition* is that modal verbs and modal auxiliaries introduce potentially unrealized states of affairs. He adds that "conferentiality by definition entails that (the NPs) have the same referents..." and therefore "...the conditions under which two conferential NPs can be identified must be the same." (p. 287). Using the terminology adopted here, one could say that if an anaphoric expression is used in situation *s*, the referent of that expression is presupposed to exist in a situation *s'* on which *s* depends. Because the real world is used as the basis for interpreting unrealized situations, conference between an anaphoric expression in the scope of a modal operator and an antecedent whose existence in the real world is presupposed is allowed. The opposite, however, is not in general the case, because the abstract situation characterizing the world does not depend on unrealized situations for its interpretation

Roberts's proposal in (Roberts, 1987; Roberts, 1989) can be seen as a way to implement Jackendoff's proposal. Roberts argues for an approach to modal subordination based on accommodation, and proposes to implement accommodation by copying DRSs. She assumes that (6.50a) and (6.50b) have incompatible interpretations of the sort in (6.55a) and (6.55b). The appropriate interpretation of (6.50b) is obtained, according to Roberts, by copying the nuclear scope of (6.55a) in the restriction of (6.55b), obtaining (6.55c), in which the discourse markers *x* and *y* are accessible to the anaphoric expression *v*. This achieves the result of what Jackendoff calls making the anaphoric expression and its antecedent depend on the same type modal operators:




The problem with Roberts' proposal is that she doesn't state the conditions under which accommodation is possible, nor how exactly it is done (i.e., what parts of a DRS are copied). In [Pesveto and



Zucchi, 1992], we note that clearly the lack of an anaphoric antecedent is not sufficient to justify accommodation, or else it would always be possible to 'subordinate' an anaphoric expression and we wouldn't have infelicitous cases such as (6.54a). We also note that the availability of telescoping the phenomenon of anaphoric reference to quantifiers that as a case of 'impossible' anaphoric reference is in many respects similar to modal subordination, is greatly facilitated whenever it is possible to interpret both the sentence containing the anaphoric antecedent and the sentence containing the anaphor as steps of single course of action (in [Poesio and Zucchi, 1992] we used the term *script*). Thus, although the continuation of (6.56a) is pretty bad for most speakers, (6.56b), in which there is a clear sense of a ceremony being described, is good, and if we add to (6.56a) the information that is required to interpret each sentence as the description of a step in the course of action, as in (6.56c), the sentence becomes much more acceptable. This suggests that the current discourse topic (what is taken to be the current discourse topic, and how easy it is to relate a new sentence to it) affects the copying process discussed by Roberts.

- (6.56) a Every student came in 79He sat near the table.  
 b Every student walked to the stage. He took his diploma, shook hands with the dean and left.  
 c The exam went on like this. Every student came in. He sat near the table, answered the professor's questions, got his mark and left.

We can hypothesize that what happens in (6.50) is as follows. As the result of the final inference step described above, the result of (6.50a) is that a new situation kind,  $k_1$ , is added to the common ground (cfr. (6.52)). This situation kind is the discourse topic of (6.50a). When processing the next utterance, i.e. (6.50b), we again have to identify the modal base of a modal, i.e., we have again to find a situation kind with respect to which to interpret the modal statement in (6.50b). The discourse topic of the previous utterance is a natural choice for such a modal base: we thus obtain the interpretation summarized in (6.57) (that I have made on purpose as close as possible to the representation used by Roberts by ignoring all details about substitutions, etc.)



since every situation that is an instance of the situation kind  $k_1$  is such that the DRS in (6.58) is verified at that situation, we can equivalently rephrase (6.57) as in (6.60), which makes the antecedent accessible to the definite description



What I have just shown is that by reformulating the interpretation of modals proposed by Kratzer as one in which it is required that a parameter ranging over situation kinds be identified, as discussed in §5.9, we do not need to assume a 'copying' operation over DRSs to account for the form of accommodation proposed by Roberts; this form of accommodation can be reformulated as a case of anaphoric interpretation, just like van der Sandt's 'syntactic' proposal about presupposition accommodation was reduced to parameter identification in §3.6

### 6.6.5 Modal Interpretation and Scope Preferences

Let's recapitulate the facts about the scopability preferences of modals discussed in §1.4. Modals tend to take wide scope with respect to indefinites, and narrow scope with respect to definites. This pattern is illustrated in (6.61a). The preference for the modal to take wide scope with respect to indefinites is not affected by the syntactic position of the indefinite, as shown in (6.61b)

- (6.61) a We should hook up an engine to the boxcar.  
 b An engine should get to Corning to pickup the boxcar.

I have analyzed in the previous section an example of interpretation that involves both modals and definite descriptions. In the current version of the theory, the scopability preferences of modals are explained as follows. When translating a sentence like (6.61a), both the discourse interpretation rules for definites and those for modals apply. Because no indefinites in our dialogues are interpreted specifically, indefinites are only ever interpreted when argument selection applies, that is, after definite interpretation and modal interpretation take place. No ambiguity is generated by the interaction of definite interpretation and modal interpretation because the result of definite interpretation makes the choice of the referent for the definite independent from the rest of the situation description, and viceversa the choice of a modal base makes the interpretation of the modal independent from the interpretation of the definite.

This account depends on letting modal interpretation taking precedence over the given/new identification procedure. The pattern in (6.61b) is also explained by the fact that both the discourse interpretation rules for definites and those for modals take precedence over the weak defaults, so that given/new partitioning takes place after the interpretation of modals has occurred

## 6.7 TENSE INTERPRETATION

As discussed in §5.6, the task of tense interpretation is to identify the course of action that includes the episode described by the utterance. Once this is done, the model construction rules for tense may apply. Further inferences lead then to the identification of the temporal relation that holds between the current episode and its preceding episode. I won't discuss these inferences here.<sup>10</sup>


### 6.7.1 The Model Construction Rule for Tense

The model construction rule for tense (both PRESENT and PAST). MCR TENSE: is a schema triggered by the occurrence inside a situation description of a maximal projection of type TP whose head is the semantic translation of tense discussed in §5.6, i.e., the function

$$\lambda \phi \lambda e [e \text{ BEFORE } \text{now} / \wedge / \text{LAST-IN}(e) \text{ ORIENTS } e] \wedge [e \in \sim / \wedge / e' = e] \wedge e \in \{\phi\}$$

MCR TENSE can be used whenever an anchor for the parameter  $e$  has been identified. The schema adds to the situation description in which the trigger occurs a new discourse marker  $e$  adds three new conditions, and replaces the logical form that triggered the schema with a description of the event  $e$  that consists of a simplified logical form from which the semantic translation of tense has been eliminated.

<sup>10</sup> A version of Dowty's TMAP could be assumed, for example, formulated as shown in [Dowty *et al.* 1997].




### 6.7.2 Identification of the Course of Action


As in Kamayama *et al.* [1993], I propose there are two kinds of interpretation rules that may assign an anchor to the situational parameter that represents the course of action of a tense operator. The first kind of rules assigns values on the basis of structural factors. The second rule generates hypotheses about the course of action on the basis of commonsense knowledge, compiled into lexical priming patterns. The structural rules are weak rules, that operate if no other rules have intervened, while the rules based on lexical priming are strong defaults and therefore take precedence over the structural rules in case of a conflict.

The rule schema MCR TENSE-STRUCTURAL is triggered by the occurrence of a TP with a non-null head, and it is located by the occurrence of the same tense operator in the same syntactic position in the logical form of the previous utterance. It should be easy to see that this condition could be given an explicit formulation with the form of common ground assumed here, it would be a matter of requiring the occurrence of an utterance prior to the current one, whose propositional content includes an underspecified representation with a parallel structure for reasons of space, however, I summarized all of these conditions into one:  $\text{POS}(e) = \text{POS}(e')$

where  $PWS$  is a function that specifies the surface structure position of an operator, and returns the path from the root of the logical form to the position of the operator, defined as the list of heads of TPs and CPs from the root of the logical form to the semantic operator. Included in the position of the parameter  $s$  matches with the position of another situational term  $s$  if they occupy the same structural position in a logical form and if all the heads of temporally relevant maximal projections (such as TP and CP) along the path also match. The position can, of course, be computed by means of structures similar to Hwang and Schubert's tense trees



I also assume that a second rule schema exists, MCR TENSE-LPRIM, with the same trigger as MCR TENSE-STRUCTURAL, but activated whenever the semantic translation of the head of the VP is lexically primed by some element of another clause. The rule also results in adding to the common ground an anchor for the parameter  $s$  that occurs in the interpretation of tense.



### 6.7.3 Tense Interpretation and Scope Disambiguation

I will first of all note that the two readings of (6.62)—a reading in which each of the members of a set of individuals perform an action, and a reading in which there is a single event characterized by the members of that set performing the action—are equivalent as long as the predicate is distributive and it can be assumed that the single actions were 'close enough' in time as to be considered a single action.

(6.62) Every participant to the manifestation was arrested.

Thus, the only 'real' cases of scope ambiguity are those in which it's possible to distinguish between a reading in which a single event is described, possibly involving all members of a set performing a distinct action, but it at approximately the same time; and a reading in which the members of the set performed the described actions in clearly distinct circumstances. An example of truly ambiguous sentences, as far as the interaction of tense with other quantifiers is concerned, is shown in (6.63). This sentence has two possible interpretations: one in which all living presidents of the US, at different times (perhaps during their mandate), shook hands with Arafat; and one in which a single event is described, such as the signature of the peace accord between Israel and the PLO that took place at the White House on Monday, September 13th, 1993.

(6.63) Every living president of the US shook hands with Arafat

When we look at the interaction of the model of tense interpretation just presented with the rest of the discourse interpretation processes discussed in this chapter, the prediction is that tense operator should 'take wide scope' over other operators whenever the choice of a course of action is motivated by world knowledge and the interpretation of other operators depends on weak defaults; *viceversa*, when an operator in the sentence is interpreted with reference to the context, while tense isn't, the operator should take wide scope with respect to scope. When neither tense nor the other operators can be interpreted as referring to information in the common ground, the choice among interpretations is left to world knowledge, or else the sentence is perceived as ambiguous.

Because of the high contextual salience of some recent events, the course of action including the event described by (6 63) is clear, and therefore my preferred interpretation of that sentence is the one in which tense takes wide scope, that is, a single event is described. In (6 64a), instead, reference is made to a contextually salient set, the set of presidents of the US, but no event is particularly salient, and therefore the universal quantifier takes wide scope. In (6 64b) no contextually relevant set is available, therefore the sentence is perceived as ambiguous.<sup>11</sup>

- (6 64) a Every (living) president of the US was born in a small town  
b Every student was born in a small town

Eng, discussing the interaction between the scope of tense and the interpretation of nominals in her thesis [Eng, 1981] observes examples like those in (6 65). Here, the choice of an interpretation seems clearly determined by world knowledge. In (6 65a), the implausibility of a situation in which all rich men are simultaneously obnoxious children; in (6 65b) conventions about dating that prescribe that one cannot date different people at the same time

- (6 65) a All rich men were obnoxious children  
b Tom dated every Miss America.

In both cases, according to the theory of discourse interpretation discussed here, two interpretations are generated by weak default rules, and the most plausible interpretation is chosen.

Another prediction of the theory of tense interpretation discussed here is that tense should take narrow scope with respect to definite descriptions, at least when these are interpreted anaphorically or with reference to the visual situation. The prediction seems to be correct

- (6 66) a The boxcar left for Corning.  
b We sent the boxcar to Corning.

The one case in which tense seems to take wide scope over definite descriptions is when the latter are interpreted generically. Consider the contrast between (6 67a) and (6 67b)


- (6 67) a The pope visited Poland  
b The pope ruled over large parts of Central Italy

<sup>11</sup> In the case of (6 64a) one might argue that, in addition, it's implausible to imagine a single event in which all US presidents being born. It is not clear whether this reasoning actually plays a role in choosing an interpretation.

## 6.8 OTHER QUANTIFIERS

There aren't enough instances of quantifiers in the TRAINS dialogues to develop a theory of the interpretation of quantifiers in that context. I will not propose other rules of interpretation beyond those discussed in §6 2 and §6 3, such as MCR PRESUPPLEMENT or SUBJECT GIVEN.

The model construction rule for universal quantifiers is very similar to that proposed in [Kamp and Reyle, 1993]



## 6.9 INDEFINITES


The only interpretation processes that apply to indefinites in our dialogues are those specified by rules such as SUBJECT GIVEN, MCR INDEF AGENT, and MCR INDEF THEME. In this section I briefly discuss the model construction rule for the presuppositional reading of indefinites and the scope effects observed in our dialogues.

### 6.9.1 Model Construction Rules

There are no model construction rules for the non-presuppositional reading of indefinites; the contribution of these NPs to the common ground is computed by rules such as MCR INDEF AGENT seen in §6 1.

The model construction rule for the presuppositional reading of indefinites, MCR INDEF PRESUPP, has the same effect as the rule for interpreting indefinites in standard DRT: it adds a new discourse marker and a set of conditions to the common ground. There are two differences: the discourse marker is always added to a situation description, instead of to an arbitrary DRS; and the rule only applies when the parameter in the indefinite has been anchored.

The rule for indefinites as subjects and agents is presented below. The rule is triggered by the occurrence of an existential in subject position, with an active verb taking a agent role as the matrix verb. The rule introduces a new discourse marker  $z$  to the current situation description, adds a condition of the form  $Q(z)$ , and rewrites the trigger by replacing the indefinite with  $z$ .



## 6.9.2 Specificity and the Scope of Indefinites

Why do indefinites always tend to take narrow scope with respect to modals and other operators in the TRAINS dialogues? I have proposed in Chapter 5 that indefinites take wide scope only when they are interpreted *specifically*, in the sense of Eng [1991].

Eng and before her Jackendoff [Jackendoff, 1972] and Grienendijk and Stokhof [Grienendijk and Stokhof, 1980], all assume that specificity is a pragmatic phenomenon. These authors all argue against the idea that specificity is a matter of wide scope with respect to another operator (as proposed, for example, by Fodor [1970]). The evidence includes the fact that the specific/non-specific ambiguity can also be had when no other operators are present. For example, Grienendijk and Stokhof discuss the following example

(6 68) A picture is missing from the museum

This sentence has two distinct readings: it may be uttered to indicate that a specific picture is missing, but it may also be used in a situation in which, say, the alarm just rang, indicating that a picture has been taken away (without the speaker knowing which particular picture). This, according to Grienendijk and Stokhof, indicates that specificity is a notion distinct from scope.

I propose that the fact that indefinites in our dialogues never 'take wide scope' with respect to modals is a consequence of the fact that in our dialogues the identity of the objects referred to with indefinites—boxcars, engines, etc.—hardly ever matters, and therefore it is not plausible for the listener to assume that the speaker meant a specific boxcar, or a specific engine.

## 6.10 ADVERBIALS

Adverbials are interesting because of the way their scope gets assigned. The Cooper storage rules in §3.2 make the semantic scope of predicate modifiers in adjunct position strictly dependent on their s-structure position. The semantic of translation of adverbials does not involve any context-dependent elements. Therefore, the model construction rules (below) do not depend on preferences, but only on tense having been interpreted.

However, when more than one eventuality is introduced in a sentence (as it is the case, for example, in sentences with modal auxiliaries) ambiguities are possible. Here is an example found in the TRAINS conversations:


(6 69) We must move an engine to Coming to pickup the boxcar

This sentence has two interpretations. In one reading, 'picking up the boxcar' is within the scope of the necessity operator: this reading can be paraphrased as follows: 'It is necessary for us to pickup the boxcar, and to do so by moving an engine to Coming.' This is the reading predicted by the semantic translation rules for adverbials attached to VPs. Another reading is also available, however: in this reading, 'picking up the boxcar' is outside the scope of the modal. This reading can be paraphrased as follows: 'For the purpose of picking up the boxcar, it is necessary for us to move an engine to Coming.'


The second reading is available, I propose, because, as discussed in §5.7, adverbials may occupy either a verb-phrase or a sentential position. The reading in which the adverbial takes narrow scope with respect to the modal is the translation of the s-structure in which the adverbial modifies the VP. The other reading is the translation of the s-structure in which the adverbial modifies the TP. In other words, I propose that (6 69) has to semantically distinct readings because it may be related to two distinct s-structures.

### 6.10.1 Model Construction Rules for Adverbials

There are two model construction rules for adverbials in VP position. The first of these is called MCR ADVF.VP. It depends on tense having been interpreted, and applies whenever the adverbial specifies a property of the eventuality (in Hwang and Schubert's terminology, whenever the adverbial is 'event-type'). The rule adds to the current situation description the property of the eventuality  $n$ , and rewrites the trigger by eliminating the adverbial.



The second rule, MCR ADVA VP, eliminates the adverbial by replacing the verb translation with a new function resulting from applying the adverbial's translation to the predicate




The rules for interpreting sentence level adverbials are similar.



## 7 Implementation: The Scope and Deindexir-3 Module of TRAINS-93

The ideas about scope disambiguation and its relation with discourse interpretation proposed in this dissertation have been implemented in a system called SAD-93. SAD-93 is a (*surface*) *disambiguate course interpretation* system given the result of syntactic interpretation and lexical disambiguation as its input. It generates a set of alternative hypotheses about the intended interpretation of a natural language utterance in a given context, for use by the Dialogue Manager component of a discourse understanding system. In doing this SAD-93 performs, in addition to scope interpretation, interpretive processes such as reference resolution and tense interpretation, the decision of integrating scope interpretation with other aspects of semantic interpretation reflects the thesis defended throughout this dissertation that the scope assigned to operators is the result of the interaction of various discourse interpretation processes.



**Figure 7.1: The architecture of TRAINS-93**

## 7.1 A GENERAL OVERVIEW OF SAD-93


SAD-93 is currently used as a module of the TRAINS-93 discourse understanding system. The TRAINS systems are intelligent planning assistants, whose domain is transportation planning

### 7.1.1 The Role of SAD-93 within TRAINS-93

The task of TRANS-93 is to assist its user in formulating a plan. The task of formulating the plan is naturally left to the user; the main task of TRANS-93 is to keep a record of the current status of the plan and try to incorporate into the plan the user's proposals. The system reports whenever an inconsistency is found and answers the user's questions. The system also has the task of verifying that the plan actually achieves the desired goal.

The architecture of TRANS-93 is schematically illustrated in Fig. 7.1. The input to TRANS-93 is first processed by the parser, a module that performs lexical disambiguation and syntactic analysis, and produces an *Indexed and Unannotated Logical Form (IULF)*, an underspecified interpretation that includes information about the lexical semantics of the words and the structure of the sentence, but no information about reference and about the scope of the operators. This representation has the form of a phrase structure tree, with the nodes labeled with the





The output of SAD 93 is a set of hypotheses about the intentions expressed by the user in uttering the sentence; each hypothesis is also represented as a SAD-expression.<sup>4</sup> The output of an utterance (7.1) is shown in (7.4).

'SAD 91 only generates hypotheses about 'surface level' speaker intentions, such as the intended interpretation of definite descriptions, tense, modals and pronouns. In general, however, these hypotheses cannot be generated in isolation from further hypotheses about the deeper intentions of the speaker as represented by speech acts and their relations (Grosz and Sidner 1986, Hobbs 1979). For this reason, the scope and denotating modules and the speech act interpretation modules will interact and will eventually be merged.

```

(7.4) ((sit-descr
      .coal
      (-dm (-col)
        (-ev-descr col
          (-drs ()
            (i user .tell :sys
              :embedding-sit
              :
              (drs (el)
                (-ev-descr el
                  (drs ()
                    (cdrs
                     (i .embedding-sit subelit s*))
                     must-occur
                     (drs (el)
                       (i el subelit s*))
                     : el has to be an instance of the
                     : kind of situation in which an
                     : action of sys and user making QJ
                     : occurs.
                     (-i el R
                       (k
                        (:lambda .e
                          (sit-descr e
                            (drs (ela)
                              : the sit includes an event of
                              : making QJ
                              (ev-descr ela
                                (drs ()
                                  (i (set-of .sys user))
                                    make :QJ))))))))))
                        (i el :at-about -col)
                        (i el subelit :embedding-sit))))))
                    (i col :at-about now))
                    (-i col :subelit coal)
                  )))
      ))

```

### 7.1.3 The Architecture of SAD-93

The architecture of SAD-93 is summarized in fig. 7.2. SAD-93 processes its input in two phases. First, *Conversational Event Generation Rules* are applied to the RIULF, yielding a preliminary description of the *Conversational Event(s)* associated with the sentence. Only the mood operators ( $\rightarrow$  DECL,  $\rightarrow$  QUES,  $\rightarrow$  TWPER) are processed in this stage. The application of the CEGK rules is a matter of straightforward pattern-matching.

Next, the phase of hypothesis propagation and generation takes place. Hypothesis generation is implemented as a forward propagation process: the higher-priority hypothesis is selected,

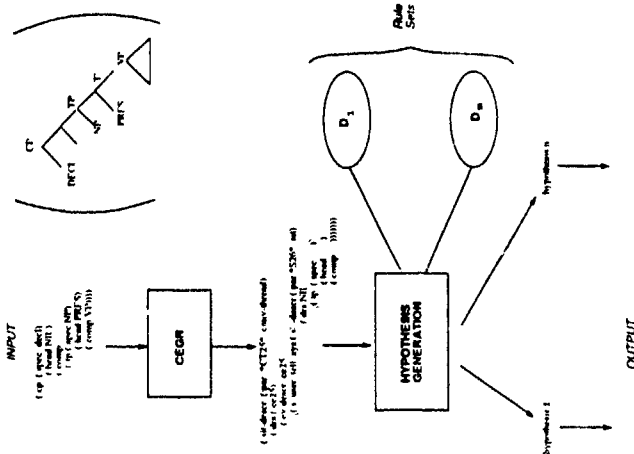


Figure 7.2 The Architecture of SAD-93

- 1 USER: We have to make OV.
- 2 There are oranges at I, and an OJ factory at B.
- 3 Engine E3 is scheduled to arrive at I at 3pm.
- 4 Shall we ship the oranges?
- 5 SYS: ok,
- 6 Shall I start loading the oranges into the empty car at I?
- 7 USER: Yes, and we'll have e3 pick it up.
- 8 OK?
- 9 SYS: OK.

Figure 7.3 The dialog processed by TRAINS-90, TRAINS-91 and TRAINS-92

then a new sub-context of the current context is created in which the content of this hypothesis is assumed. Finally, all applicable rules are applied to generate new hypotheses, that are put on the agenda. This phase is started by taking the result of the Conversational Event Generation Rules as the initial hypothesis. The rules used to generate new hypotheses encode the discourse interpretation and DRS rewriting axioms described in Chapter 6. Hypothesis generation is described in § 3.1.1. § 3.2.1 discusses the form of rules used in SAD-93 and how they are applied.

The precedence relations between defaults are implemented by dividing the rules used to generate new hypothesis sets in databases according to their priority. During the hypothesis propagation phase, the system applies first the rules in the highest priority database, that encode the higher priority discourse interpretation processes such as definite description interpretation. Then the system applies the rules in the rule database that contains the rules with second highest priority, as well as all the rules in the first database. These rules encode weaker defaults, such as those based on the thematic hierarchy. The rules in the databases at higher priority may prevent the rules in the lower priority databases from applying by rewriting the logical form.

The hypothesis processing method implemented in SAD-93 incorporates steps to eliminate duplicate hypotheses, thus, the presence of alternative hypotheses at the end of the hypothesis generation phase means that, as far as SAD-93 is concerned, the input was truly ambiguous in the given context. The task of dealing with the ambiguity is left to the dialogue manager and the plan recognition module, if one of the hypotheses is more plausible, that one is chosen, otherwise, a repair sequence may be initiated.

#### 7.1.4 The Dialogues Handled by the System

The phenomena that TRAINS-93 is designed to handle are illustrated by the two dialogues in Fig. 7.3 and Fig. 7.4.

Each of these dialogues is the edited transcription of a spoken conversation between two humans playing one the role of system, the other the role of user. The experimental setting is as follows<sup>3</sup>. The two humans are in the same room, separated by a partition, and they speak via

<sup>3</sup>A different setting later abandoned, was used to collect the dialog in Fig. 7.3

- 1 USER: we better ship a boxcar of oranges to Bath by 8 AM.  
 2 SYS: okay.  
 3-3-6 USER: so we need to get a boxcar to Corning  
 where there are oranges.  
 3-7 USER: there are oranges at Corning.  
 3 8 SYS: right?  
 4 SYS: right.  
 5-1 USER: so we need an engine to move the boxcar.  
 5-2 USER: right?  
 6 SYS: right.  
 7-1-2 USER: so there's an engine at Avon.  
 7-3 USER: right?  
 8 SYS: right.  
 9-13 USER: so we should move the engine at Avon, engine E1,  
 to Dansville to pick up the boxcar there.  
 14 SYS: okay.  
 15-2-4 USER: and move it from Dansville to Corning.  
 15-5-7 load up some oranges into the boxcar.  
 15-8-10 USER: end then move it on to Bath.  
 16 SYS: okay.  
 17 USER: how does that sound?  
 18-3 SYS: That's no problem.  
 19 USER: good.

Figure 7.4 The dialog processed by TRAINS-93




Figure 7.5 The map used by the participants in the conversation.

microphones and headphones. They have in front of them a map of the TRAINS domain—a map containing a list of towns, together with the resources at every town. (The map used by the participants in Fig. 7.4 is shown in Fig. 7.5.) The human playing the role of user learns of the transportation goal to be achieved by randomly choosing a card on which the task is reported.

The operators to be scripted in these dialogues include tense, definite and indefinite descriptions, noun, modalis, and wh-phrases. The referential expressions include definite descriptions and pronouns.

## 7.2 THE SYNTAX OF SAD-EXPRESSIONS

The BNF definition of the syntax of SAD-expressions is shown in Fig. 7.6 to 7.8. SAD-expressions are Lisp lists, whose CAR indicates the type of expression (All atoms are keywords). The main kinds of SAD-expressions are

1. SAD-expressions representing constituents of syntactic trees (maximal projections and adjunctive structures): e.g.,  
`(:vp (:head (:sem :move)) (:comp (:np (:sem :Eng-1))))`  
 These SAD-expressions may be used as conditions in DRSs
2. SAD-expressions representing DRSs: these are lists whose first element is the keyword `:drs`, followed by a list of markers, followed by zero or more conditions, as in  
`(:drs (:x ..... :y) (:i :x :p) (:i :y :Q))`
3. SAD-Expressions representing situation descriptions and event descriptions. These are of the form `(:sit-descr :s (:drs .....))` and `(:ev-descr :e (:drs .....))`, respectively
4. SAD-Expressions representing Episodic Logic expressions. These SAD-expressions are used to represent the logical semantics of words; infixed expressions may also be used as conditions of DRSs

```

<adj-expression> → (s-structure) | (vt-descr)
<s-structure> → (maximal projection) | (adjoining structure)
(maximal projection) → ( (maximal projection type)
  ((type (subtree)))
  (head (subtree))
  ((comp (subtree)))
)
(maximal projection type) → (:cp | :tp | :vp | :vp' | :op | :pp
(subtree) → (s-structure) | (lexical semantics expression)
(lexical semantics expression) → (item (SEL-expression))
(adjoining structure) →
  ( (adjoining structure type) (adjoined (subtree))
  (adjoined (subtree))) |
  ( (adjoining structure type) (adjoined (subtree))
  (adjoined (subtree)))
(adjoining structure type) → :np-adjunct | :cp-adjunct | :vp-adv |
:vp-adv'

```

Figure 7.6. BNF Definition of the Syntax of SAD-expressions (1)

```

(vt-descr) → (:vt-descr (symbol) (DRS))
(vt-descr) → (:vt-descr (symbol) (DRS))
'DRS' → (:drs (markers) (condition)*)
(markers) → NIL | ( (symbol) (markers) )
(condition) → (s-structure) | (complex condition) |
  (:vt-descr) | (ev-descr) | (logical expression) |
  (infix expression) | (marker specification)
(complex condition) → (:cnds (DRS) (operator) (DRS))
(marker specification) → (:dm (symbol))
(SEL-expression) → (infix expression) | (quantified expression) |
  (modal expression) | (logical expression)
(infix expression) → (:i (term) (infix operator) (term)*)
(quantified expression) → (:q (quantifier) (SEL-expression))
(modal expression) → (:m (symbol) (lambda expression)
  (lambda expression))

```

Figure 7.7. BNF Definition of the Syntax of SAD-expressions (2)

```

(logical expression) → (:= (term) (term)) |
  (=anchor (parameter) (term))
(infix operator) → (predicate) | (and | or | :^ | :R | :subid)
(term) → (Lisp keyword) | (kind) | (indexical) | (parameter) | :s^
(predicate) → (Lisp keyword) | (functional expression) |
  (:pair (symbol)) | (lambda expression)
(functional expression) → (:f (predicate) (term) *)
(lambda expression) → (:lambda (symbol) (SEL expression))
(quantifier) → (:op (determiner) (symbol) (SEL expression))
(kind) → (:k (predicate))
(indexical) → (:indexical (symbol))
(parameter) → (:pair (symbol) (parameter type))
(parameter type) → :we | :it | :he | :I | :all | :shared-sit | :conv-thread
(determiner) → :e | :the | :every

```

Figure 7.8 BNF Definition of the Syntax of SAD-expressions (3)




Figure 7.9 Generation of new hypothesis trees.

## 7.3 DETAILS ABOUT THE IMPLEMENTATION

### 7.3.1 Hypothesis Generation

SAD-93 generates hypotheses about the intended interpretation of an utterance by forward propagation. This process is realized as a straightforward breadth-first search implemented with a queue. Whenever the queue is not empty, the next hypothesis tree is extracted. (The hypothesis tree is the data structure used to represent an hypothesis. An hypothesis, in general, does not consist of a single fact, but of a collection of them, the hypothesis tree stores pointers to all the facts that are part of an hypothesis.)

Once an hypothesis tree has been selected, a context switch takes place, to a new context in which that hypothesis is assumed: the context inherits all the facts in the current context and, in addition, includes the facts in the selected hypothesis tree. (The use of contexts is discussed in §7.3.3.) All the facts in the selected hypothesis tree are then considered, to see if any of them triggers a rule among those in the current rule database. Whenever a rule is activated, the operators specified by that rule are performed, resulting in the creation of a new hypothesis tree, as shown in Fig. 7.9; rule application and rule definition are discussed in §7.3.2. Once all of the facts in the hypothesis tree have been considered, the hypothesis tree is put in the closed queue.

Each rule is included in one of several databases (currently, five) organized by priority. The rules that are always active (e.g., the model construction rules) go in the database with priority 0. The rules with immediately lower priority go in the database of priority 1, which also inherits all the rules in the database of priority 0; and so forth. Hypothesis generation begins by first applying only the rules in the database of priority 1 are processed, this database contains most of the discourse interpretation rules, as well as all the rules in the database at priority 0. When

no more rules in this database can apply, the resulting hypothesis trees are obtained from the closed queue, these are the hypothesis trees from which no further hypothesis tree was obtained. The hypothesis trees that result from phase 1 are put back in the queue, the database containing the rules with priority 2, 1 and 0 becomes active, and the hypothesis propagation process is started again. At the end of this second phase, all the resulting hypothesis trees are enqueued again and the database with priority 3 becomes active, which contains all rules at level 0, 1 and 2.

The hypothesis propagation algorithm does not use an evaluation function to order the hypothesis trees. The hypothesis trees are enqueued in a first-come, first-serve fashion. The system does include a pruning mechanism, however, to eliminate those branches of the search tree that could only lead to duplicate hypotheses. This is quite crucial, because a list of duplicate hypothesis trees may be generated by e.g., applying in different permutations two rules that do not affect each other's outcome (such as two rules anchoring two distinct parameters).

In order to recognize duplicate hypothesis trees, the hypothesis tree data structure has a rule history field that records the 'path' from the initial hypothesis tree to this hypothesis tree: the rules have been applied together with the patterns that triggered them. A hash table is also used, indexed by (rule pattern) pairs, the value of the hash table at a pair  $r, p$  is the set of hypothesis trees whose history includes an application of rule  $r$  to pattern  $p$  at some point. This hash table is updated whenever a new hypothesis tree is obtained. When an hypothesis tree  $H$  is extracted from the queue, the system checks first of all if  $H$  is a duplicate of an hypothesis already processed, it does so by checking if another hypothesis tree  $H'$  is in all entries of the hash table indexed by one of the elements in the rule history of  $H$ , that is, if there is another hypothesis tree obtained from the initial hypothesis by the same rules that resulted in the generation of  $H$ , although perhaps in a different order. An example of this kind of check is shown in Fig. 7.10 if HT  $k$  has been obtained by applying RULE 2 to the same pattern, which resulted in HT 1 HT  $k$  will be pruned because the subtree rooted at HT  $k$  would not contain any hypothesis trees included in the subtree rooted at HT 1.

The main algorithms used in hypothesis propagation are shown in Algol-like format in Fig. 7.11.




Figure 7.10. Pruning hypothesis trees.

### 7.3.2 Rules

The rules used in SAD-93 are of a very simple sort. A rule specification consists of three parts:

- a trigger to be matched,
- a set of constraints to be verified
- a set of events that must occur whenever the trigger matches and the constraints are verified.

The rule processing method in SAD-93 is also largely based on what is done in standard forward propagation algorithms. There are, however, additional complexities both in the rule specification language and in the rule processing method that derive from the fact that these rules are meant to implement the DRS construction algorithm.

First of all, we need both rules whose trigger has to unify with a complete fact and rules that are activated whenever their trigger unifies with an expression occurring anywhere in the SAD-expression representing the fact. (From now on, I'll use the term *fact expression* for this kind of expressions.) Rules of the first kind are used to encode logical axioms and/or the sort of axioms that encode commonsense reasoning, most of the rules formalizing DRS rewriting are, however, of the second kind. For example, the DRS construction rule for universal NPs (discussed later in this section) is triggered by the occurrence anywhere in the RULF of a universal NP of the form in (7.5). In fact, at the current stage of development of SAD-93 there are many more rules of the second than of the first type.

(7.5)  $((np \text{ issem } \lambda p \forall x [x \models Q(x)] P(x)))$



```

PROPAGATE (INPUT I)
BEGIN
  hypothesis-trees-to-propagate ← (CREATE-HYPOTHESIS-TREE(input-I))
  FOR rule-db IN rule-databases DO
    rule-databases ← a list of
    rule-databases ordered by priority
  BEGIN
    open-queue ← NIL, closed-queue ← NIL
    FOR ht IN hypothesis-trees-to-propagate DO
      ENQUEUE(open-queue)
    CLOSE(rule-db)
    hypothesis-trees-to-propagate ← NIL
    FOR ht IN closed-queue DO
      IF FORMAL-HYPOTHESIS-TEST(ht) THEN
        hypothesis-trees-to-propagate ←
          hypothesis-trees-to-propagate UNION {ht}
      END
    END
    hypothesis-trees-to-propagate ←
      hypothesis-trees-to-propagate UNION {ht}
  END
  RETURN hypothesis-trees-to-propagate
END

CLOSE(rule-db)
RETURN
WHILE current-hypothesis-queue ≠ DEQUEUE(open-queue) DO
  BEGIN
    IF "REFUNDANT"(current-hypothesis-tree) THEN
      BEGIN
        FOR fact IN current-hypothesis-tree DO
          FOR ht IN rule-db DO
            APPLY-RULE(rule-fact) DO
              ENQUEUE(open-queue)
            END
          END
        ENQUEUE(current-hypothesis-tree, closed-queue)
      END
    END
  END
END

```

Figure 7.11 Hypothesis propagation algorithms

Furthermore, the trigger matching algorithm must be relatively flexible. In particular, it must be capable of detecting partial matching. For example, a trigger representing a syntactic subphrase of the form  $(:xp (:head ?EXP))$  must match a SAD-expression of the form  $(:spec <EXP1> (:head <EXP2>))$ ; most importantly, a trigger representing a DRS, such as  $(:drs (?x) (:i ?x :P))$ , must match a SAD-expression including additional markers and/or conditions, such as

( $:drs (:x :y) (:i :x :P) (:i :y :Q)$ )

In general, what is required is for a rule to become active on a fact expression whenever its trigger subsumes the fact expression, or one of its subexpressions. Subsumption between SAD-expressions is defined as follows

**Definition 7.1** The SAD-expression  $se$  subsumes the SAD-expression  $se'$  if either (i)  $se$  is a pattern variable, or (ii) if they are expressions of the same type (i.e., both infix expressions, both DRS, etc.) and furthermore

- 1 If both expressions are Lisp keywords they are identical, else,
- 2 If both expressions are either terms, predicates, SEL-expressions, marker specifications, lexical semantics expressions or situation or event descriptions, they have the same number of elements and each element of  $se$  subsumes the corresponding elements of  $se'$ , else,
- 3 If both expressions are DRSs, each marker in  $se$  subsumes a marker in  $se'$ , and each condition in  $se$  subsumes a condition in  $se'$ , else,
- 4 If both expressions are s-structures, each existing constituent of  $se$  subsumes the corresponding constituent in  $se'$  (e.g., the head of  $se$  subsumes the head of  $se'$ , etc.).

In SAD-93, a fact triggers a rule if the trigger of the rule subsumes the fact in the sense just specified; and the rule definition language is such that the default behavior of the system is to recursively attempt matching a rule's trigger with each subexpression of a fact expression, the pattern has to be explicitly marked in order to prevent the recursion, as discussed below.

Additional complications are involved in specifying what a rule can/must do once it becomes active. In a traditional rule-based system, there are two basic operations a rule can perform: to add new facts, and to delete existing facts. In SAD-93 deletion is not really needed, at least for the moment, but we do need the capability to rewrite the current fact, this is what most of the DRS construction rules do.<sup>4</sup> In addition, the SAD-expression to be rewritten is often embedded in the fact expression, typically, but not always, the SAD-expression to be rewritten is the one matched by the trigger (see however below). Therefore, the semantics of the rewrite operation has been defined so that by default, the expression that gets rewritten is the expression subsumed

<sup>4</sup>Of course, a more general solution would be obtained by having a system able to deal with addition and deletion, and by implementing rewrite in terms of the two. For the moment, only the simpler system has been developed.

```

(defrule PCR-PRES
  (trigger (tp
    (spec TNP)
    (head
      (sem
        (lambda (p)
          (lambda (x)
            (op e ?ev (lambda (y)
              (lambda (z)
                (at-about ?pos) and
                (lambda (w)
                  (lambda (v)
                    (comp ?NP))))))))
            (comp ?NP))))
          (constraints (global (anchored ?sem nil)))
          (events
            (add (dm ?ev))
            (add (lambda (z)
              (lambda (w)
                (at-about ?pos))
              (rewrite (ev-decor ?ev (dis (lambda (t)
                (spec TNP)
                (head NIL)
                (comp ?NP))))))
            ))
  ))

```

Figure 7.12 The model construction rule for present tense

by the trigger. The semantics of the *add* operation is defined so that, by default, the expression is added to the DRS that contains the subexpression being triggered.

As an example, consider the definition of the model construction rule for the present tense. This rule is triggered by the occurrence within the fact form of a TP whose head is the following function

```

(lambda p
  (lambda x
    (lambda (y)
      (lambda (z)
        (lambda (w)
          (lambda (v)
            (op e ?ev (lambda (y)
              (lambda (z)
                (at-about (indexical ?now)) and
                (lambda (w)
                  (lambda (v)
                    (comp ?NP))))))))
            (comp ?NP))))
          (constraints (global (anchored ?sem nil)))
          (events
            (add (dm ?ev))
            (add (lambda (z)
              (lambda (w)
                (at-about ?pos))
              (rewrite (ev-decor ?ev (dis (lambda (t)
                (spec TNP)
                (head NIL)
                (comp ?NP))))))
            ))
  ))

```

the rule becomes active in a context in which the parameter  $(\text{par } :S^* :s\text{it})$  is anchored (which requires the prior activation of a discourse interpretation rule). The DRS construction rule introduces an event description and locates the event at the *now* point. That DRS construction rule is implemented in SAD-93 by the rule in Fig. 7.12. The trigger of the rule is an expression representing a maximal projection of type *tp*; this maximal projection may occur anywhere in the fact form. The events that follow the activation of the rule include (i) adding to the DRS containing the triggering sub-expression a discourse marker (this is what the sub-expression  $(\text{dm } ?ev)$  does, more on this below) (ii) adding to the same DRS a condition of the form  $(\text{lambda } ?ev : \text{at-about } ?pos)$ , and (iii) rewriting the triggering expression by deleting the head of the maximal projection. I will say more on the meaning of *global* shortly.

More formally, the syntax of rule definition is as follows.

```

(7.6) (rule def) →
      (Defrule (rule name)
        (trigger (pattern))
        ((constraints (constraint)+))
        ((events (event)+)))

```

The trigger of a rule is a single pattern. The syntax of patterns is the same as the syntax of SAD-expressions defined in §7.2, except that (i) a pattern may be used in place of any simple component of a sub-expression, and (ii) a pattern may be of the form  $(\text{lambda } ?x \text{ } \langle \text{pattern} \rangle)$ . A variable is a symbol prefixed by a question mark, such as  $?x$  or  $?y$ ;  $?x$  stands for 'Root DRS' and may only appear in a pattern as the first element of the top list?

If the trigger of a rule is of the form  $(\text{lambda } ?x \text{ } \langle \text{pattern} \rangle)$ , the trigger matches a SAD-expression if its second element subsumes the SAD-expression. Otherwise, the trigger matches a fact-expression if it subsumes any subexpression in the fact.

A rule is activated on a fact if its trigger matches it and, in addition, its set of constraints is verified. The syntax for specifying the constraints of a rule is shown in (7.7).

```

(7.7) (constraints) → (pattern) |
      // (pattern) |
      (lisp (lisp-function) (arg)+ (result))) |
      :global (pattern))

```

A constraint is either a pattern, possibly including variables, or a call to a Lisp function. The set of constraints is used as a query to the facts database (see §7.3.3); all possible ways of satisfying the conjunction of constraints in the current context are found. The set of constraints is verified if there is at least one substitution of values for the variables in the pattern that results in a goal that can be proved. Each distinct substitution results in a distinct rule application, hence a distinct hypothesis tree being created. If a constraint is of the form  $(\text{lambda } ?x \text{ } \langle \text{pattern} \rangle)$ , the goal is proved in the root DRS. The form  $(\text{lambda } ?x \text{ } \langle \text{pattern} \rangle)$  is used to evaluate some special forms (including  $:=$ ,  $:=$  anchor and  $:=$  anchored) in the outermost database. As explained in §7.3.3, the backward rewriter that processes the queries to the facts database deals with an expression of the form

```
(:lisp (function name) <args> <result>))
```

by replacing the variable bindings in  $\langle \text{args} \rangle$ , invoking on the arguments the Lisp function  $\langle \text{function name} \rangle$ , and unifying the result with  $\langle \text{result} \rangle$ .

The syntax for specifying the events to occur when a rule is activated is given in (7.8).

```

(7.8) (events) → (lisp (lisp-function) (arg)+ (result)))
      (add (pattern))
      (rewrite-when (pattern) (pattern))

```

<sup>1</sup>Currently, there is an additional constraint on patterns: they must be equivalent to their normalized form, which means they cannot include more than one condition per DRS. See below.

```

(defrule MCR-EVERY
  (trigger (tp
    (sem (op every ?x (1 ?res-elt * (1 ?x ?01))))
    , to be verified LOCALLY (as part of
    , the same situation descriptions of
    , which the trigger is a part))
  (constraints (global (anchored ?res-elt)))
  (events (create-dm 8 ?x1))
  (1 ?tp (create-dm 8 ?x2))
  (rewrite ?x)
  (rewrite-when
    (tp (spec ?RUBN)
      (head ?HEAD)
      (comp ?COMP))
    (cdrs
      (dm ?x1)
      (sit-descr ?x1 (dm ?x1) (1 ?x ?01))
      (1 ?x1 subel ?x1))
    every
    (dm ?x2)
    (sit-descr ?x2
      (dm 1 (tp
        (spec ?RUBN)
        (head ?HEAD)
        (comp ?COMP)))
      (1 ?x2 subel ?x1))))
  , a universal condition
  )

```

Figure 7.13 The model construction rule for universals

As said above, each add event adds the SAD-expression resulting from replacing variables with their bindings either to the DRS which contains the triggering expression, or to the root DRS if the pattern is of the form  $(/ \langle \text{expression} \rangle)$ . A rewrite event rewrites the expression subsumed by the trigger. Rewrite-when events are used when more than one rewrite is needed. For example, the rule for universal quantification involves two rewrites. The NP that triggers the rule is rewritten and replaced in the logical form by a discourse marker; in addition, the first embedding TP must be replaced by a complex condition with operator *every* that includes the TP in which the NP has been replaced by a discourse marker in its nuclear scope. The first pattern in the specification of the event gives the context in which the rewrite has to occur, the second pattern indicates the SAD-expression that results from the rewrite. The way in which the rule *MCR-EVERY* makes use of rewrite-when events is shown in Fig. 7.13.

Trigger matching is the most common operation performed by the rule processing component of SAD-93, it is therefore essential that it be efficient. One way to implement the matching algorithm could be to visit the whole fact form recursively for each rule, attempting to determine for each subexpression of the fact form whether the trigger of the rule subsumes that subexpression. This strategy, however, involves duplicating lots of work. For this reason, the trigger matching algorithm actually visits the fact form only once, for each subexpression, it finds all

rules whose trigger subsumes the current subexpression, tries to verify their constraints, and eventually executes the operations.

One must also be careful in computing DRS subsumption: in general, determining whether a DRS  $K_1$  subsumes a DRS  $K_2$  requires proving, for each condition in  $K_1$ , whether that condition is included in  $K_2$ , which is as complex as proving a goal in a database. A direct implementation of DRS subsumption along these lines proved too slow. Looking at the rule triggers, however, one can see that none of them involves more than one condition per DRS. This observation suggested the following implementation. All facts are normalized before being added to an hypothesis tree. Normalization involves (1) splitting each DRS so that no more than one condition occurs in each form of type  $(\text{discs } \dots)$ , and (2) representing the membership of a discourse marker  $X$  in a DRS by adding to the DRS a condition of the form  $(\text{dm } X)$  instead of adding the discourse marker in the list in second position of a DRS form. Normalization of a DRS produces a set of normalized DRSs, as shown in example (7.9). Subsumption among normalized DRSs can be computed by simple unification. Of course, DRS normalization generates a large number of facts to be considered by the rule application algorithm.

```

(7.9) (discs (x :y) (1 :x :p) (1 :y :R :x))
      =>
      (discs n11 (dm :x))
      (discs n11 (dm :y))
      (discs n11 (1 :x :x :p))
      (discs n11 (1 :x :f :R :x))

```

The definition of the *APPLY-RULE* function, as well as the modified version of *CLOSURE* that takes care of visiting the fact form recursively, are presented in Fig. 7.14. *MDB:ASK* is the interface to the simple backward reasoner described in §7.3.3.

The interface with the rule databases is provided by the following two macros

- `(RYP :: IN-RULE-DB <RULE-DB-NAME>)` can be used to change the database to which rules are added.
- `(RYP :: EXPRULEDB <ruledb> (:SUPER <rule-do>+))` can be used to define a rule database and the database it inherits rules from

### 7.3.3 The Fact Database

A simple database package has been developed, whose basic building blocks—unification and querying functions—derive from those in the Prolog interpreter in Lisp developed in [Norvig, 1992]. The unifier actually served as basis for all variable-binding operations in SAD-93. This package has been intentionally designed as a module separate from the rest of SAD-93, so that it could easily be replaced by a more sophisticated package.

The facts database is partitioned into contexts, hierarchically organized. Associated with each context are a list of 'local' axioms, together with a list of superior contexts from which facts are inherited. The fact database can be augmented by means of the function *MDB:TELL*, and can be queried by means of the function *MDB:ASK*. Both functions take a list of axioms

```

(CLOSURE (rule-db)
  BEGIN
    WHILE current-hypothesis-tree = DEQUEUE () DO
      BEGIN
        IF = REDUNDANT?(current-hypothesis-tree) THEN
          BEGIN
            MDS ASSUME FACTS-IN (current-hypothesis-tree)
            BEGIN
              FOR fact IN current-hypothesis-tree DO
                BEGIN
                  FOR rule IN rule-db DO
                    FOR ht IN APPLY-RULE(rule, fact) DO
                      ENQUEUE(ht, open-queue)
                    FOR subexp IN SUBEXPRESSIONS(fact) DO
                      FOR ht IN APPLY-RULE(rule, subexp) DO
                        ENQUEUE(ht, open-queue)
                      END
                    END
                  ENQUEUE(current-hypothesis-tree, closed-queue)
                END
              END
            END
          END
        APPLY-RULE (rule, fact)
        BEGIN
          hts = NIL
          FOR m IN SUBEXPRESSIONS(fact) DO
            FOR b IN MDS ASK (rule constraints (current-db)) DO
              hts = UNION hts b
            END
          END
          hts = EVALUATE (COMPUTE-EVENTS-RESULT(rule, events, hts))
        END
      END
    END
  END
)

```

Figure 7.14 The rule application algorithm

and optionally a context as their arguments. The macro MDS:ASSUME can be used to query the database in a context that includes an additional set of axioms

The database package includes a simple form of equality reasoning, extensively used in SAD-93 for anchoring parameters. Equality assertions are made relative to contexts, and can be inherited just like all other assertions, for example, two terms  $X$  and  $Y$  can be asserted to be equal by calling MDS:TELL as follows:

```
(MDS:TELL ' ((:= :X :Y)) <db>)
```

The unification algorithm in Norvig's book has been modified to work with the equality subsystem, so that a goal of the form  $(:= :X :P)$  can be satisfied in a database in which the equality assertion above is stored together with the fact  $(:= :Y :P)$

The querying function MDS:ASK has been modified so that calls to Lisp functions can occur in the list of goals. These calls can be used to bind variables to the result of a call to a Lisp function, for example,

```
(.LISP (SUBST 'C 'D '(A B D) 'PRES))
```

binds  $\text{PRES}$  to  $(A B C)$

### 7.3.4 Discourse Markers and Priming

SAD-93 maintains two hash tables: one, indexed by discourse marker, records the type of each situation marker and the situation that it describes; the other, indexed by situation identifiers, records the discourse markers that occur in each situation. These tables are used for reference resolution purposes and also to compute lexical priming relations used for interpreting anaphoric definite descriptions and tense. At the moment, only NPs can be primed, and only by other NPs with the same type. For example, the NP "a boxcar," used in sentence 1 will prime the NP "the boxcar" used in sentence 2, but the NP "a house" does not, in the current version of the system, prime the NP "the door," as in (7.10a) nor does the verb "eat" prime the NP "the food," as in (7.10b).

- (7.10) a Harvey walked towards a house. THE DOOR was open  
b Kim ate out yesterday. THE FOOD was good

## 8 Conclusions

The main idea presented in this dissertation is the proposal that the scope preferences observed in the literature are not the result of an independent 'scope disambiguation' process. Instead, they are the result of independent interpretation procedures, none of which especially concerned with 'scope disambiguation,' all working off underspecified interpretations. I have provided a formal characterization of discourse interpretation in terms of parameter anchoring and identified a class of sentence constituents called *operators* whose interpretation is context-dependent and whose scope is not determined by s-structure. I have also discussed in some detail several interpretation procedures: definite description interpretation, the interpretation of modals, the identification of the given/new partition, argument selection, and tense interpretation. Along the way, I proposed a model of discourse interpretation called *Conversational Representation Theory* that builds upon Discourse Representation Theory but incorporates the major hypotheses about ambiguity and discourse interpretation that I proposed in the dissertation.

The primary motivation for this work has been the intuition, backed by all the psychological evidence I am aware of, that people have preferred interpretations for the kind of 'scopally ambiguous' sentences one finds in naturally occurring conversations (i.e., leaving aside artificially created examples, about which people may indeed have no intuitions). This is especially true, it seems, when the subjects are involved in a task such as planning orange shipments or answering questions. The second motivation was the observation, also resulting from the psychological literature, that none of the disambiguation factors proposed in the literature could explain all preferences, yet there was no clear proposal about which factors did play a role and how they interacted.

But even more compelling, I'd say, was the feeling that the existing proposals about assigning a scope to operators were all oblivious to the fact that scope is in fact a semantic notion, and were essentially based on operations of logical form manipulation, augmented with various heuristics. This lack of a good understanding of the process by which semantic scope is assigned was also reflected by the fact that most of the existing work on scope was concerned with making sure that all semantically available readings were generated.

## Bibliography

- [Atney, 1987] Atney, S. P. 1987. *The English Noun Phrase in its Sentential Aspect*. Ph.D. Dissertation, MIT, Cambridge MA.
- [Akmajian et al., 1979] Akmajian, A.; Steele, S.; and Wasow, T. 1979. The category AUX in universal grammar. *Linguistic Inquiry* 10:1-64.
- [Allen and Perrault, 1980] Allen, J. F. and Perrault, C. 1980. Analyzing intention in utterances. *Artificial Intelligence* 15(3):143-178.
- [Allen and Schubert, 1991] Allen, J. F. and Schubert, L.K. 1991. The TRAINS project. TRAINS Technical Note 91-1, University of Rochester, Department of Computer Science.
- [Allen et al., 1990] Allen, J. F.; Hendler, J.; and Tate, A., editors. 1990. *Readings in Planning*. Morgan Kaufmann, San Mateo, CA.
- [Allen, 1987] Allen, J. F. 1987. *Natural Language Understanding*. Benjamin Cummings, Menlo Park, CA.
- [Allport, 1987] Allport, A. 1987. Visual attention. In Posner, M. I., editor. 1987. *An Introduction to Cognitive Science*. The MIT Press, Cambridge, MA. chapter 16, 631-682.
- [Alshawi and Crouch, 1992] Alshawi, H. and Crouch, R. 1992. Monotonic semantic interpretation. In *Proc. 30th. ACL*, University of Delaware. 32-39.
- [Alshawi, 1992] Alshawi, H., editor. 1992. *The Core Language Engine*. The MIT Press.
- [Altman and Steedman, 1983] Altman, G. T. M. and Steedman, M. 1988. Interaction with context during human sentence processing. *Cognition* 30:191-238.
- [Altman, 1989] Altman, G. T. M., editor. 1989. *Parsing and Interpretation*. Lawrence Erlbaum, Hove, East Sussex, UK.
- [Appelt and Konolige, 1988] Appelt, D. and Konolige, K. 1988. A practical nonmonotonic theory of reasoning about speech acts. In *Proc. ACL-88*, Buffalo.
- [Austin, 1962] Austin, J. L. 1962. *How to Do Things with Words*. Harvard University Press, Cambridge, MA.

- [Bach, 1981] Bach, E. 1981. On time, tense, and aspect. An essay in English metaphysics. In Cole, P., editor 1981, *Radical Pragmatics*. Academic Press, New York.
- [Barwise and Cooper, 1981] Barwise, J. and Cooper, R. 1981. Generalized quantifiers and natural language. *Linguistics and Philosophy* 4(2) 159-219.
- [Barwise and Cooper, 1993] Barwise, J. and Cooper, R. 1993. Extended Kamp notation. In Azbel, P., Israeli, D., Kagan, Y., and Peters, S., editors 1993, *Situation Theory and its Applications*, v.3. CSLI chapter 2, 29-54.
- [Barwise and Perry, 1983] Barwise, J. and Perry, J. 1983. *Situations and Attitudes*. The MIT Press.
- [Barwise, 1987] Barwise, J. 1987. Noun phrases, generalized quantifiers and anaphora. In Gärdenfors [1987] 1-30.
- [Barwise, 1989] Barwise, J. 1989. *The Situation in Logic*. CSLI Lecture Notes. University of Chicago Press.
- [Belnap, 1982] Belnap, N. 1982. Questions and answers in Montague grammar. In Peters, S. and Saarinen E., editors 1982, *Processes, Beliefs and Questions*. D. Reidel, Dordrecht.
- [Bennett and Partee, 1978] Bennett, M. and Partee, B. 1978. *Toward the Logic of focus and aspect in English*. Indiana University Linguistics Club, Bloomington, IN.
- [Berman, 1987] Berman, S. 1987. Situation-based semantics for adverbs of quantification. In *Proceedings of WCCFL 6*. Stanford Linguistic Association.
- [Berman, 1991] Berman, S. 1991. *On the Semantics and Logical Form of wh-clauses*. Ph.D. Dissertation, University of Massachusetts at Amherst.
- [Bresnan, 1970] Bresnan, J. 1970. On complementizers: toward a syntactic theory of complement types. *Foundations of Language* 6 297-321.
- [Brewka, 1991] Brewka, G. 1991. *Nonmonotonic Reasoning*. Logical Foundations of Computation. Cambridge University Press.
- [Bunt, 1985] Bunt, H. C. 1985. *Max Terms and Model Theoretic Semantics*. D. Reidel.
- [Carberry 1990] Carberry, S. 1990. *Plan Recognition in Natural Language Dialogue*. The MIT Press, Cambridge, MA.
- [Carlson and Tannenhaus, 1988] Carlson, G. N. and Tannenhaus M. 1988. Thematic roles and language comprehension. In Wilkins, W., editor 1988, *Syntax and Semantics 21: Thematic Relations*. Academic Press, New York. 263-288.
- [Carlson, 1978] Carlson, G. N. 1978. *Reference in Kinds in English*. Ph.D. Dissertation, University of Massachusetts, Amherst.
- [Carlson, 1984] Carlson, G. N. 1984. On the role of thematic roles in linguistic theory. *Linguistics* 22 259-279.

- [Carter, 1987] Carter, D. M. 1987. *Interpreting Anaphors in Natural Language Texts*. Ellis Horwood, Chichester, UK.
- [Charniak and Goldman, 1988] Charniak, E. and Goldman, R. P. 1988. A logic for semantic interpretations. In *Proc. ACL 88*, Buffalo, NY, 87-94.
- [Charniak, 1986] Charniak, E. 1986. A new theory of marker passing. In *Proceedings of the Fifth National Conference on Artificial Intelligence*, Philadelphia, Pennsylvania. American Association for Artificial Intelligence. 584-588.
- [Chierchia and McConnell-Ginet, 1990] Chierchia, G. and McConnell-Ginet, S. 1990. *Meaning and Grammar: An Introduction to Semantics*. The MIT Press, Cambridge, MA.
- [Chierchia and Turner, 1988] Chierchia, G. and Turner, R. 1988. Semantics and property theory. *Linguistics and Philosophy* 11:261-302.
- [Chierchia, 1984] Chierchia, G. 1984. *Topics in the Syntax and Semantics of Infinitives and Gerunds*. Ph.D. Dissertation, University of Massachusetts at Amherst.
- [Chierchia, 1992] Chierchia, G. 1992. Anaphora and dynamic logic. *Linguistic and Philosophy* 15: 111-183.
- [Chierchia, 1993] Chierchia, G. 1993. Questions with quantifiers. *Natural Language Semantics* 1: 181-234.
- [Chomsky, 1957] Chomsky, N. 1957. *Syntactic Structures*. Mouton, The Hague.
- [Chomsky, 1967] Chomsky, N. 1967. *Aspects of the Theory of Syntax*. MIT Press. Cambridge, MA.
- [Chomsky, 1981] Chomsky, N. 1981. *Lectures on Government and Binding*. Paris, Dordrecht.
- [Chomsky, 1982] Chomsky, N. 1982. *Some Concepts and Consequences of the Theory of Government and Binding*. The MIT Press.
- [Chomsky, 1985a] Chomsky, N. 1985a. *Barriers*. MIT Press, Cambridge, MA.
- [Chomsky, 1985b] Chomsky, N. 1985b. *Knowledge of Language*. Praeger Publications, New York.
- [Clark and Brennan, 1990] Clark, H. H. and Brennan, S. E. 1990. Grounding in communication. In Resnick, L. B.; Levine, J., and Behrend, S. D., editors 1990, *Perspectives on Socially Shared Cognition*. APA.
- [Clark and Haviland, 1977] Clark, H. H. and Haviland, S. E. 1977. Comprehension and the given-new contract. In Freedle, R. O., editor 1977, *Discourse Production and Comprehension*. Ablex, Norwood, NJ, chapter 1, 1-40.
- [Clark and Marshall, 1981] Clark, H. H. and Marshall, C. R. 1981. Definite reference and mutual knowledge. In *Elements of Discourse Understanding*. Cambridge University Press, New York.

- [Clark and Schaefer, 1989] Clark, H. H. and Schaefer, E. F. 1989. Contributing to discourse. *Cognitive Science* 13 259-94.
- [Clark, 1993] Clark, H. H. 1993. *Aspects of Language Use*. University of Chicago Press.
- [Cocchiarella, 1976] Cocchiarella, N. 1976. On the logic of natural kinds. *Philosophy of Science* 43 202-222.
- [Cohen and Levesque, 1990] Cohen, P. R. and Levesque, H. J. 1990. Rational interaction as the basis for communication. In Cohen, P. R., Morgan, J. and Pollack, M., editors 1990, *Intentions in Communication*. Morgan Kaufmann, chapter 12, 221-256.
- [Cohen and Perrault, 1979] Cohen, P. R. and Perrault, C. R. 1979. Elements of a plan based theory of speech acts. *Cognitive Science* 3(3) 177-212.
- [Cohen, 1978] Cohen, P. R. 1978. On knowing what to say. Planning speech acts. Technical Report 118, Department of Computer Science, University of Toronto.
- [Cohen, 1984] Cohen, P. R. 1984. The pragmatics of referring and the mobility of common action. *Computational Linguistics* 10(2) 97-146.
- [Cooper 1983] Cooper, R. 1983. *Quantification and Situation Theory*. D. Reidel Publishing Company, Dordrecht, Holland.
- [Cooper 1993] Cooper, R. 1993. Generalised quantifiers and resource situations. In Aczel, P. Israel, D., Katagiri, Y. and Peters, S., editors 1993, *Situation Theory and its Applications*, v. 1. CSLI and University of Chicago, Stanford, chapter 8, 191-212.
- [Crain and Steedman, 1985] Crain, S. and Steedman, M. 1985. On not being led up the garden path: the use of context by the psychological syntax processor. In Dowty, D. R., Karttunen, L. and Zwicky, A. M., editors 1985, *Natural Language Parsing: Psychological, Computational and Theoretical Perspectives*. Cambridge University Press, New York, 320-358.
- [Cresswell, 1973] Cresswell, M. J. 1973. *Logics and Languages*. Methuen and Co., London.
- [Dalrymple et al., 1991] Dalrymple, M.; Shieber, S. M.; and Pereira, F. C. N. 1991. Ellipsis and higher-order unification. *Linguistics and Philosophy* 14(4) 399-452.
- [de Jong, 1987] Jong, E. de 1987. The compositional nature of (in)definiteness. In Reuland, E. J. and Meulen, A. G. B. ter, editors 1987, *The Representation of (In)definiteness*. The MIT Press, Cambridge, MA, chapter 11, 270-285.
- [de Swart, 1991] Swart, H. de 1991. *Adverbs of quantification: a generalized quantifier approach*. Ph.D. Dissertation, Rijksuniversiteit Groningen.
- [de Swart, 1992] Swart, H. de 1992. Intervention effects, monotonicity and scope. In Barker, C. and Dowty, D., editors 1992, *SALT II: Proceedings of the Second Conference on Semantics and Linguistic Theory*, Columbus, OH: Ohio State University, 387-406.
- [Dekker and Hendriks, 1993] Dekker, P. and Hendriks, H. 1993. Situation semantics. Unpublished manuscript.

- [Delia, 1989] Delia, Judy 1989. *Cleft Constructions*. Ph.D. Dissertation, University of Edinburgh, Edinburgh, Scotland.
- [Devlin, 1991] Devlin, K. 1991. *Logic and Information*. Cambridge University Press, Cambridge, UK.
- [Di Eugenio, 1993] Di Eugenio, B. 1993. *Understanding Natural Language Instructions: A Computational Approach to Purpose Clauses*. Ph.D. Dissertation, University of Pennsylvania, Philadelphia.
- [Diesing, 1992] Diesing, M. 1992. *Indefinites*. The MIT Press, Cambridge, MA.
- [Dowty et al., 1981] Dowty, D. R.; Wall, R. E.; and Peters, S. 1981. *Introduction to Montague Semantics*. D. Reidel, Dordrecht, Holland.
- [Dowty, 1979] Dowty, D. R. 1979. *Word Meaning and Montague Grammar*. D. Reidel, Dordrecht, Holland.
- [Dowty, 1986a] Dowty, D. R. 1986a. The effects of aspectual class on the temporal structure of discourse: Semantics or pragmatics? *Linguistics and Philosophy* 9(1).
- [Dowty, 1986b] Dowty, D. R. 1986b. A note on collective predicates, distributive predicates, and all\*. In *Proc. ESCOL-86*.
- [Dowty, 1989] Dowty, D. R. 1989. On the semantic content of the notion of 'thematic role'. In Partee, B., Chierchia, G. and Turner, R., editors 1989, *Properties, Type, and Meaning*, volume 2. Kluwer, Dordrecht, 69-130.
- [Dowty, 1991] Dowty, D. R. 1991. Thematic proto-roles and argument selection. *Language* 67(3) 547-619.
- [E- 1981] Eng, M. 1981. *Tense Without Scope: An Analysis of Nouua as Indexicals*. Ph.D. Dissertation, University of Wisconsin, Madison, WI. Distributed by IULC.
- [Eng, 1991] Eng, M. 1991. The semantics of specificity. *Linguistic Inquiry* 22(1) 1-25.
- [Engdahl, 1986] Engdahl, E. 1986. *Constituents Questions*. D. Reidel, Dordrecht.
- [Fagin and Halpern, 1988] Fagin, R. and Halpern, J. Y. 1988. Belief, awareness, and limited reasoning. *Artificial Intelligence* 34:39-76.
- [Feastad et al., 1987] Feastad, J.E.; Halvorsen, P.K.; Langholm, T.; and Benihem, J. van 1987. *Situations, Language and Logic*. D. Reidel, Dordrecht.
- [Ferguson and Allen, 1993] Ferguson, G. and Allen, J. F. 1993. Cooperative plan reasoning for dialogue systems. In *Proceedings of the AAAI-93 Fall Symposium on Collaborative Human-Computer Interaction: Reconciling theory, synthesizing practice*, Raleigh, NC.
- [Ferguson, 1992] Ferguson, G. M. 1992. Explicit representation of events, actions, and plans for assumption-based plan reasoning. Technical Report 428, University of Rochester, Rochester, NY.

- [Fikes and Nilsson, 1971] Fikes, R. E. and Nilsson, N. J. 1971. STRIPS: A new approach to the application of theorem proving to problem solving. *Artificial Intelligence* 2:189-208.
- [Fillmore, 1968] Fillmore, C. 1968. The case for case. In Bach, C. and Harms, R., editors 1968, *Universals in Linguistic Theory*. Holt, Rinehart and Winston, New York. 1-81.
- [Fodor and Sag, 1982] Fodor, J. D. and Sag, I. A. 1982. Referential and quantificational indices. *Linguistics and Philosophy* 5:355-398.
- [Fodor 1970] Fodor, J. D. 1970. *The Linguistic Description of Opaque Contexts*. Ph.D. Dissertation, MIT, Cambridge, MA.
- [Fodor, 1982] Fodor, J. D. 1982. The mental representation of quantifiers. In Peters, S. and Saarinen, E., editors 1982, *Processors: Beliefs and Questions*. D. Reidel, Dordrecht.
- [Fox, 1987] Fox, B. A. 1987. *Discourse Structure and Anaphora*. Cambridge University Press, Cambridge, UK.
- [Frazier and Fodor, 1978] Frazier, L. and Fodor, J. D. 1978. The sausage machine: A new two stage parsing model. *Cognition* 6:291-295.
- [Gabbay and Reyle, 1986] Gabbay, D. and Reyle, U. 1986. Direct deductive computation on Discourse Representation Structures: part I. Deduction using run time skolemization. To Appear.
- [Gärdenfors, 1987] Gärdenfors, P., editor 1987. *Generalized Quantifiers*. D. Reidel, Dordrecht, The Netherlands.
- [Gärdenfors, 1988] Gärdenfors, P. 1988. *Knowledge in Flux: Modeling the Dynamic of Epistemic States*. The MIT Press, Cambridge, MA.
- [Gärdenfors, 1992] Gärdenfors, P., editor 1992. *Belief Revision*. Cambridge University Press, Cambridge.
- [Gawron and Peters, 1990] Gawron, J. M. and Peters, S. 1990. *Anaphora and Quantification in Situation Semantics*, volume 19 of *Lecture Notes*. CSLI.
- [Gazdar et al., 1985] Gazdar, G.; Klein, E.; Pullum, G.; and Sag, I. 1985. *Generalized Phrase Structure Grammar*. Blackwell, Oxford.
- [Gazdar, 1979] Gazdar, G. 1979. *Pragmatics*. Academic Press, New York.
- [Ginzburg, 1992] Ginzburg, J. 1992. *Questions, Queries and Facts: a Semantics and Pragmatics for Interrogatives*. Ph.D. Dissertation, Stanford University.
- [Givón, 1989] Givón, T. 1989. *Mind Code and Context*. Lawrence Erlbaum.
- [Givón, 1980] Givón, D. S., editor 1989. *Resolving Semantic Ambiguity*. Springer Berlin.
- [Grice, 1967] Grice, H. P. 1967. Logic and conversation. (William James Lectures.) In *Studies in the Way of Words*, Cambridge. Harvard University Press, 1989.

- [Grice, 1969] Grice, H. P. 1969. Vague names. In Davidson, D. and Hintikka, J., editors 1969, *Words and Objections*. Reidel, Dordrecht. 118-145.
- [Grimshaw, 1990] Grimshaw, J. 1990. *Argument Structure*. The MIT Press.
- [Groenendijk and Stokhof, 1990] Groenendijk, J. A. G. and Stokhof, M. J. B. 1990. A pragmatic analysis of specificity. In Heny, F., editor 1990, *Ambiguities in Intentional Context*. D. Reidel, Dordrecht and Boston. 153-190.
- [Groenendijk and Stokhof, 1989] Groenendijk, J. A. G. and Stokhof, M. J. B. 1989. Type shifting and the semantics of interrogatives. In Chierchia, G., Partee, B. and Turner, R., editors 1989, *Properties Type and Meaning*, v.2. D. Reidel, Dordrecht.
- [Groenendijk and Stokhof, 1991] Groenendijk, J. A. G. and Stokhof, M. J. B. 1991. Dynamic Predicate Logic. *Linguistics and Philosophy* 14:39-100.
- [Gross et al., 1993] Gross, D., Allen, J., and Traum, D. 1993. The TRAINS 91 dialogues. TRAINS Technical Note 92-1, Computer Science Dept. University of Rochester.
- [Grosz and Sidner, 1986] Grosz, B. J. and Sidner, C. L. 1986. Attention, intention, and the structure of discourse. *Computational Linguistics* 12(3):175-204.
- [Grosz et al., 1986] Grosz, B. J., Jones, K., Sparck, and Webber, B. Lynn editors 1986. *Readings in Natural Language Processing*. Morgan Kaufmann, Los Altos, CA.
- [Grosz et al., 1987] Grosz, B. J., Appelt, D., Martin, P., and Pereira, F. 1987. TEAM: An experiment in the design of transportable natural language interfaces. *Artificial Intelligence* 32(2):173-244.
- [Grosz, 1977] Grosz, B. J. 1977. *The Representation and Use of Focus in Dialogue Understanding*. Ph.D. Dissertation, Stanford University.
- [Gruber, 1965] Gruber, J. S. 1965. *Studies in Lexical Relations*. Ph.D. Dissertation, MIT Cambridge, MA. Distributed by Indiana University Linguistics Club.
- [Gundel, 1974] Gundel, J. K. 1974. *The Role of Topic and Comment in Linguistic Theory*. Ph.D. Dissertation, University of Texas at Austin. Reprinted by Carl and Publishing, New York and London, 1988.
- [Haegeman, 1991] Haegeman, L. 1991. *An Introduction to Government and Binding Theory*. Basil Blackwell, first edition.
- [Hamblin, 1987] Hamblin, C. 1987. *operatives*. Basil Blackwell.
- [Harman, 1977] Harman, G. 1977. *Review of linguistic behavior* by Jonathan Bennett. *Language* 53:417-424.
- [Haviland and Clark, 1974] Haviland, S. E. and Clark, H. H. 1974. What's new? acquiring new information as a process in comprehension. *Journal of Verbal Learning and Verbal Behavior* 13:512-521.



- [Hawkins, 1978] Hawkins, J. A. 1978. *Definiteness and Indefiniteness*. Croom Helm, London.
- [Hayes-Roth, 1985] Hayes-Roth, F. 1985. Rule-based systems. *Communications of the ACM* 28(9) 921-932.
- [Hayes, 1987] Hayes, Patrick J. 1987. What the frame problem is and isn't. In Polyzdym, Z. W., editor 1987, *The Robot's Dilemma: The Frame Problem in Artificial Intelligence*. Theoretical Issues in Cognitive Science. Ablex, Norwood, New Jersey. 123-137.
- [Heenan, 1993] Heenan, P. A. 1993. Speech actions and mental states in task-oriented dialogues. In *Working Nines AAAI Spring Symposium on Reasoning about Mental States: Formal Theories & Applications*. Stanford.
- [Heim, 1982] Heim, I. 1982. *The Semantics of Definite and Indefinite Noun Phrases*. Ph D. Dissertation, University of Massachusetts at Amherst.
- [Heim, 1983a] Heim, I. 1983a. File change semantics and the familiarity theory of definiteness. In Bauerle, R.; Schwarze, C.; and Stechow, A. von, editors 1983a. *Meaning Use and Interpretation of Language* de Gruyter, Berlin.
- [Heim, 1983b] Heim, I. 1983b. On the projection problem for presuppositions. In *Proc WCCFL II*, Stanford, CA.
- [Heim, 1987] Heim, I. 1987. Where does the definiteness restriction apply? In Meulen, A. ter and Reuland, E., editors 1987, *The Representation of (In)definiteness*. MIT Press.
- [Heim 1990] Heim, I. 1990. E-type pronouns and donkey anaphora. *Linguistics and Philosophy* 13(2) 137-158.
- [Hendrix, 1978] Hendrix, G. G. 1978. Semantic aspects of translation. In Walker, D., editor 1978, *Understanding Spoken Language*. Elsevier, New York. 193-226.
- [Herzog and Rollinger, 1991] Herzog, O. and Rollinger, C.-R. 1991. *Text understanding in LILOG*. Springer Verlag, Berlin and New York.
- [Hindle and Rooth, 1993] Hindle, D. and Rooth, M. 1993. Structural ambiguity and lexical relations. *Computational Linguistics* 19(1) 103-120. Special Issue on Using Large Corpora.
- [Hinnrichs, 1981] Hinnrichs, E. 1981. *Temporale anaphora in englischen StaatsExamen thesis*, Universitt Tbingen.
- [Hinnikka, 1974] Hinnikka, J. 1974. Quantifiers vs quantificational theory. *Linguistic Inquiry* 5 153-177.
- [Hirst, 1987] Hirst, G. 1987. *Semantic Interpretation and the Resolution of Ambiguity*. Studies in Natural Language Processing. Cambridge University Press, Cambridge, UK.
- [Hirst, 1990] Hirst, G. 1990. Mixed-depth representation for natural language text. In *Proc. AAAI Spring Symposium* 10-14.

- [Hitzman, 1993] Hitzman, J. M. 1993. *Temporal Adverbials and the Syntax-Semantics Interface*. Ph D. Dissertation, University of Rochester, Department of Linguistics, Rochester, NY. Also available as TR 931, Cognitive Science Program, University of Rochester.
- [Hobbs and Martin, 1987] Hobbs, J. R. and Martin, P. 1987. Local pragmatics. In *Proc. IJCAI-87*, Milano, Italy. 520-523.
- [Hobbs and Shieber, 1987] Hobbs, J. R. and Shieber, S. M. 1987. An algorithm for generating quantifier scopeings. *Computational Linguistics* 13(1-2) 47-63.
- [Hobbs et al., 1990] Hobbs, J. R.; Stickel, M.; Martin, P.; and Edwards, D. 1990. Interpretation as abduction. Technical Note 499, SRI International, Menlo Park, CA.
- [Hobbs, 1979] Hobbs, J. R. 1979. Coherence and coreference. *Cognitive Science* 3 67-90.
- [Hobbs, 1982] Hobbs, J. R. 1982. Representing ambiguity. In *Proc WCCFL* 15-28.
- [Hobbs, 1983] Hobbs, J. R. 1983. An improper treatment of quantification in ordinary English. In *Proc ACL-83*, Cambridge, MA. 57-63.
- [Horn, 1989] Horn, Laurence R. 1989. *A Natural History of Negation*. The University of Chicago Press.
- [Hudson, 1988] Hudson, S. B. 1988. *The Structure of Discourse and Anaphor Resolution: The Discourse Center and the Roles of Nouns and Pronouns*. Ph D. Dissertation, University of Rochester.
- [Hurum and Schubert, 1986] Hurum, S. and Schubert, L. K. 1986. Two types of quantifier scoping. In *Proceedings Sixth CSCSI/SCEIO*, 39-43.
- [Hurum, 1988] Hurum, S. 1988. Handling scope ambiguities using domain-independent heuristics. Technical Report TR 88-12, University of Alberta, Edmonton, Canada.
- [Hwang and Schubert, 1992] Hwang, C. H. and Schubert, L. K. 1992. Tense trees as the "fine structure" of discourse. In *Proc. ACL-92*, Newark, DE. 232-240.
- [Hwang and Schubert, 1993] Hwang, C. H. and Schubert, L. K. 1993. Episodic logic: A situational logic for natural language processing. In Azeai, P.; Israel, D.; Katagiri, Y.; and Piers, S., editors 1993, *Situation Theory and its Applications*, v.3. CSLI. 303-338.
- [Hwang, 1992] Hwang, C. H. 1992. *A Logical Approach to Narrative Understanding*. Ph D. Dissertation, University of Alberta, Department of Computing Science, Edmonton, Alberta, CA.
- [Ioup, 1975] Ioup, G. 1975. Some universals for quantifier scope. In Kimball, J., editor 1975, *Syntax and Semantics 4*. Academic Press, New York. 37-58.
- [Jackendoff, 1972] Jackendoff, R. 1972. *Semantic Interpretation in Generative Grammar*. MIT Press, Cambridge, MA.
- [Jackendoff, 1983] Jackendoff, R. 1983. *Semantics and Cognition*. The MIT Press.

- [Jackendoff, 1987] Jackendoff, R. 1987. The status of thematic relations in linguistic theory. *Linguistic Inquiry* 18, 369-411.
- [Jackendoff, 1990] Jackendoff, R. 1990. *Semantic Structures*. The MIT Press, Cambridge, MA.
- [Johnson and Klein, 1986] Johnson, M. and Klein, E. 1986. Discourse, anaphora and parsing. Technical Report CSLI-86-43, CSLI, Stanford, CA.
- [Jones, 1985] Jones, C. F. 1985. *Syntax and Semantics of Infinitival Adjuncts*. Ph.D. Dissertation, University of Massachusetts at Amherst, Linguistics, Amherst, MA.
- [Kadmon, 1987] Kadmon, N. 1987. *On Unique and Non-Unique Reference and Asymmetric Quantification*. Ph.D. Dissertation, University of Massachusetts at Amherst.
- [Kálmán, 1990] Kálmán, L. 1990. Deferred information: The semantics of commitment. In Kálmán, L. and Pócs, L., editors 1990, *Papers from the Second Symposium on Logic and Language*. Akadémiai Kiadó, Budapest, 125-148.
- [Kameyama et al., 1993] Kameyama, M., Pasouneau, R.; and Poesio, M. 1993. Temporal centering. In *Proc. ACL-93*, Columbus, OH.
- [Kamp and Reyle, 1991] Kamp, H. and Reyle, U. 1991. A calculus for first order Discourse Representation Structures. Sonderforschungsbereich 340 Bericht 14, University of Stuttgart, Stuttgart.
- [Kamp and Reyle, 1993] Kamp, H. and Reyle, U. 1993. *From Discourse to Logic*. D. Reidel, Dordrecht.
- [Kamp and Röddescheider, 1992] Kamp, H. and Röddescheider, A. 1992. Remarks on lexical structure, dis-construction and lexically driven inference. Arbeitspapiere des SFB 340 21, IMS, University of Stuttgart, Stuttgart, Germany.
- [Kamp, 1979] Kamp, H. 1979. Events, instant and temporal reference. In Bauerle, R., Egli, U., and Stechow, A. von, editors 1979, *Semantics from Different Points of View*. Springer-Verlag, 376-417.
- [Kamp, 1981] Kamp, H. 1981. A theory of truth and semantic representation. In Groenendijk, J., Janssen, T., and Stokhof, M., editors 1981. *Formal Methods in the Study of Language*. Mathematical Centre, Amsterdam.
- [Kamp, 1990] Kamp, H. 1990. Prolegomena to a structural account of belief and other attitudes. In Anderson, C. A. and Owens, J., editors 1990, *Propositional Attitudes—The Role of Content in Logic, Language and Mind*. University of Chicago Press and CSLI, Stanford, chapter 2, 27-50.
- [Karttunen and Peters 1979] Karttunen, L. and Peters, S. 1979. Conventional implicature. In *Syntax and Semantics 11: Presupposition*, Academic Press, New York, 1-56.
- [Karttunen 1976] Karttunen, L. 1976. Discourse referents. In McCawley, J., editor 1976, *Syntax and Semantics 7: Notes from the Linguistic Underground*. Academic Press, New York.

- [Karttunen, 1977] Karttunen, L. 1977. The syntax and semantics of questions. *Linguistics and Philosophy* 1, 1-44.
- [Katz, 1980] Katz, J. J. 1980. Chomsky on meaning. *Language* 56, 1-41.
- [Kautz, 1987] Kautz, H. A. 1987. *A Formal Theory of Plan Recognition*. Ph.D. Dissertation, Department of Computer Science, University of Rochester, Rochester, NY. Also available as TR 215, Department of Computer Science, University of Rochester.
- [Kayne, 1984] Kayne, R. 1984. *Connectedness and Binary Branching*. Foris, Dordrecht.
- [Keenan, 1976] Keenan, E. L. 1976. Towards a universal definition of "subject". In Li, C., editor 1976, *Subject and Topic*. Academic Press, New York, 303-333.
- [Keenan, 1987] Keenan, E. L. 1987. Unreducible n-ary quantifiers in natural language. In Gardent, [1987] 109-150.
- [Keller, 1988] Keller, W. R. 1988. Nested cooper storage: The proper treatment of quantification in ordinary noun phrases. In Reyle, U. and Rohrer, C., editors 1988, *Natural Language Parsing and Linguistic Theories*. D. Reidel, Dordrecht, 432-447.
- [Kempson and Cormack, 1981] Kempson, R. and Cormack, A. 1981. Ambiguity and quantification. *Linguistics and Philosophy* 4(2), 259-310.
- [Kiss, 1993] Kiss, K. E. 1993. Wh-movement and specificity. *Natural Language and Linguistic Theory*.
- [Koenig, 1993] Koenig, J. P. 1993. Mapping constructions as word templates: Evidence from French. In *Proc. of the 6th Biennial Conference on Grammatical Relations*.
- [Kratzer, 1977] Kratzer, A. 1977. What 'must' and 'can' must and can mean. *Linguistics and Philosophy* 1, 337-355.
- [Kratzer, 1981] Kratzer, A. 1981. The notional category of modality. In Eikemeier, H. J. and Rieger, H., editors 1981, *World, Worlds, and Context: New Approaches in Word Semantics*. de Gruyter, Berlin, 38-74.
- [Kratzer, 1989] Kratzer, A. 1989. An investigation of the lumps of thought. *Linguistics and Philosophy* 12, 607-653.
- [Kripke, 1972] Kripke, S. A. 1972. Naming and necessity. In Davidson, D. and Harman, G., editors 1972, *Semantics of Natural Language*. Reidel, Dordrecht, 253-355.
- [Kripke, 1977] Kripke, S. A. 1977. Speaker reference and semantic reference. In French, P. A., Uehling, T. E., and Wettstein, H. K., editors 1977, *Contemporary Perspectives in the Philosophy of Language*. University of Minnesota Press, Minneapolis, 6-27.
- [Krohn, 1975] Krohn, A. 1975. *The Semantics of Scope in English*. Ph.D. Dissertation, MIT, Cambridge, MA.

- [Kuno, 1971] Kuno, S. 1971. The position of locatives in existential sentences. *Linguistic Inquiry* 2(3).
- [Kurtzman and MacDonald, 1993] Kurtzman, H. S. and MacDonald, M. C. 1993. Resolution of quantifier scope ambiguities. *Cognition* 48 243-279.
- [Ladusaw, 1977] Ladusaw, W. 1977. *Polarity Sensitivity as inherent scope relations*. Ph.D. Dissertation, University of Texas, Austin.
- [Lakoff, 1970] Lakoff, G. P. 1970. A note on vagueness and ambiguity. *Linguistic Inquiry* 1(1) 357-359.
- [Lakoff, 1971] Lakoff, G. P. 1971. Semantic interpretation in generative grammar. In Steinberg and Jakobovits (1971).
- [Landman, 1986a] Landman, F. 1986a. Data semantics: an epistemic theory of partial objects. In *Towards a Theory of Information* (1986a) 1-96.
- [Landman, 1986b] Landman, F. 1986b. *Towards a Theory of Information*. Groningen-Amsterdam Studies in Semantics. Temis, Dordrecht.
- [Landman, 1989a] Landman, F. 1989a. Groups I. *Linguistics and Philosophy* 559-605.
- [Landman, 1989b] Landman, F. 1989b. Groups II. *Linguistics and Philosophy* 723-744.
- [Langholm, 1988] Langholm, T. 1988. *Partiality, Truth and Persistence*. CSLI Lecture Notes. University of Chicago and CSLI, Stanford, CA.
- [Larson, 1988] Larson, R. K. 1988. On the double object construction. *Linguistic Inquiry* 19(3) 335-391.
- [Lascarides and Asher, 1991] Lascarides, A. and Asher, N. 1991. Discourse relations and defeasible knowledge. In *Proc. ACL-91*, University of California at Berkeley 55-63.
- [Lascarides et al., 1992] Lascarides, A.; Asher, N.; and Oberlander, J. 1992. Inferring discourse relations in context. In *Proc. ACL-92*, University of Delaware 1-8.
- [Levesque, 1990] Levesque, H. J. 1990. All I know: A study in autopsystemic logic. *Artificial Intelligence* 42 263-309.
- [Levin, 1993] Levin, B. 1993. *English Verb Classes and Alternations*. University of Chicago Press, Chicago and London.
- [Levinson, 1983] Levinson, S. 1983. *Pragmatics*. Cambridge University Press.
- [Lewis, 1969] Lewis, D. K. 1969. *Convention: A Philosophical Study*. Harvard University Press, Cambridge, MA.
- [Lewis, 1975] Lewis, D. K. 1975. Adverbs of quantification. In Keenan, E., editor. 1975. *Formal Semantics of Natural Language*. Cambridge University Press, Cambridge.

- [Lewis, 1978] Lewis, D. K. 1978. Scorekeeping in a language game. *Journal of Philosophical Logic* 8:339-359.
- [Linde, 1978] Linde, C. 1978. Focus of attention and the choice of pronouns in discourse. In Grønn, T., editor. 1978. *Syntax and Semantics 12*. Academic Press.
- [Lusk, 1983] Lusk, G. 1983. The logical analysis of plurals and mass terms: A lattice-theoretical approach. In Bläugle, R.; Schwart, C.; and Stockow, A., eds., editors. 1983. *Meaning, Use and Interpretation of Language*. Walter de Gruyter, 302-323.
- [Lodibaz, 1987] Lodibaz, S. 1987. *Definites*. *Journal of Semantics* 4:279-326.
- [Lyons, 1977] Lyons, J. 1977. *Semantics*. Cambridge University Press.
- [Marden-Wilson and Tyler 1980] Marden-Wilson, W.D. and Tyler, L.K. 1980. The temporal structure of spoken language understanding. *Cognition* 8:1-71.
- [Marden-Wilson, 1975] Marden-Wilson, W.D. 1975. Sentence perception as an interactive parallel process. *Science* 189:226-228.
- [Marden-Wilson, 1989] Marden-Wilson, W.D., editor. 1989. *Lexical Representation and Processes*. MIT Press, Cambridge, MA.
- [May, 1977] May, R. 1977. *The Grammar of Quantification*. Ph.D. Dissertation, MIT, Cambridge, MA.
- [May, 1985] May, R. 1985. *Logical Form in Natural Language*. The MIT Press.
- [McCawley, 1988] McCawley, J. D. 1988. *The Syntactic Phenomena of English* (volumes 1 and 2). The University of Chicago Press.
- [Morris, 1991] Morris, A. 1991. Imperatives: Linguistics vs. philosophy. *Linguistics* 29(669-702).
- [Milsark, 1974] Milsark, G. 1974. *Existential Sentences in English*. Ph.D. Dissertation, MIT, Cambridge, MA.
- [Milsark, 1977] Milsark, G. 1977. Towards an explanation of certain peculiarities in the existential construction in english. *Linguistic Analysis* 3:1-30.
- [Moens and Steedman, 1988] Moens, M. and Steedman, M. 1988. Temporal ontology and temporal reference. *Computational Linguistics* 14(2):15-28.
- [Molinsman, 1992] Molinsman, F. 1992. *Coordination and Comparatives*. Ph.D. Dissertation, MIT.
- [Morris, 1988] Morris, D. 1988. Quantifier scoping in the SRL Core Language Engine. In *Proc. ACL-88*, Buffalo, NY, 33-40.
- [Moshons, 1989] Moshons, R. 1989. *Meaning and Partiality*. Ph.D. Dissertation, University of Amsterdam.

- [Nakhimovsky, 1988] Nakhimovsky, A. 1988. Aspect, aspectual class, and the temporal structure of narratives. *Computational Linguistics* 14(2) 29-43.
- [Neale, 1990] Neale, S. 1990. *Descriptions*. The MIT Press, Cambridge, MA.
- [Nebel, 1990] Nebel, Bernhard 1990. *Reasoning and Reason in Hybrid Representation Systems*. Springer-Verlag.
- [Nishidauchi, 1984] Nishidauchi, T. 1984. Control and the thematic domain. *Language* 60 215-260.
- [Norvig, 1987] Norvig, P. 1987. Inferences in text understanding. In *Proc AAAI-87*. Seattle, Washington.
- [Norvig, 1992] Norvig, P. 1992. *Paradigms of artificial intelligence programming*. Morgan Kaufman, San Mateo, CA.
- [Palmer, 1986] Palmer, F. R. 1986. *Mind and Modality*. Cambridge University Press, Cambridge, UK.
- [Parsons, 1990] Parsons, T. 1990. *Events in the Semantics of English*. The MIT Press, Cambridge, MA.
- [Partee and Routh, 1983] Partee, B. H. and Routh, M. 1983. Generalized conjunction and type ambiguity. In Balerio, R., Schwarze, C., and Stechow, A. von, editors 1983. *Meaning, Truth and Interpretation of Language*. Walter de Gruyter, Berlin, West Germany.
- [Partee, 1972] Partee, B. H. 1972. Opacity, coreference, and pronouns. In Davidson, D. and Harman, G., editors 1972. *Semantics for Natural Language*. D. Reidel, Dordrecht, Holland 415-441.
- [Partee, 1973] Partee, B. H. 1973. Some structural analogies between tenses and pronouns in English. *Journal of Philosophy* 70 601-609.
- [Partee, 1984] Partee, B. H. 1984. Nominal and temporal anaphora. *Linguistics and Philosophy* 7 243-286.
- [Pereira and Pollack, 1991] Pereira, F. C. N. and Pollack, M. E. 1991. Incremental interpretation. *Artificial Intelligence* 40 37-82.
- [Pereira, 1999] Pereira, F. C. N. 1999. A calculus for semantic composition and scop-ing. In *Proc ACL-99*.
- [Pereira, 1990] Pereira, F. C. N. 1990. Categorical semantics and scop-ing. *Computational Linguistics* 16(1) 1-10.
- [Perrault, 1990] Perrault, C. R. 1990. An application of default logic to speech act theory. In Cohen, P. R., Morgan, J., and Pollack, M. F., editors 1990. *Formalisms in Communication*. The MIT Press, Cambridge, MA. chapter 9, 161-185.

- [Piercey, 1982] Piercey, D. 1982. *Paths and Categories*. Ph.D. Dissertation, MIT, Cambridge, MA.
- [Pescio and Zucchi, 1992] Pescio, M. and Zucchi, A. 1992. On telescoping. In *Proceedings SALZ-2*, Columbus, Ohio.
- [Pescio, 1991] Pescio, M. 1991. Relational semantics and scope ambiguity. In Barwise, J., Gärden, J. M., Pridem, G.: and Tarsky, S., editors 1991. *Situation Semantics and its Applications*, vol 2. CSLI, Stanford, CA. chapter 20, 469-497.
- [Pescio, 1993] Pescio, M. 1993. A situation-theoretic formalization of definite description interpretations in plan elaboration dialogues. In Azzel, P., Israel, D., Katsagi, Y., and Peters, S., editors 1993. *Situations Theory and its Applications*, vol 3. CSLI, Stanford. chapter 12, 139-174.
- [Pescio, 1994] Pescio, M. 1994. Weak definites. In *Proceedings of the Fourth Conference on Semantics and Linguistic Theory (SALT-4)*, Cornell University Press.
- [Pudlak, 1989] Pudlak, J.-Y. 1989. Verb movement, universal grammar, and the structure of IP. *Linguistic Inquiry* 20 365-424.
- [Ryle, 1988] Ryle, David L. 1988. A logical framework for default reasoning. *Artificial Intelligence* 36(1).
- [Rysdal, 1971] Rysdal, P. V. 1971. *Cross-over Phenomena*. Holt, Rinehart and Winston, New York.
- [Rysdal, 1974] Rysdal, P. M. 1974. *On Raising*. The MIT Press, Cambridge, MA.
- [Prince, 1981] Prince, E. F. 1981. Toward a taxonomy of given-new information. In Cole, P. editor 1981. *Radical Pragmatics*. Academic Press, New York. 223-236.
- [Prince, 1988] Prince, A. N. 1988. *Past, Present and Future*. Clarendon Press, Oxford.
- [Reichman, 1985] Reichman, R. 1985. *Getting Computers to Talk Like You and Me*. The MIT Press, Cambridge, MA.
- [Reinhart, 1976] Reinhart, T. 1976. *The Semantic Domain of Anaphora*. Ph.D. Dissertation, MIT, Cambridge, MA.
- [Reinhart, 1981] Reinhart, T. 1981. Pragmatics and linguistics: An analysis of sentence topics. *Philosophica* 27(1). Also distributed by Indiana University Linguistics Club.
- [Reinhart, 1983] Reinhart, T. 1983. *Anaphora and semantic interpretation*. Croom Helm, London.
- [Rosen, 1980] Rosen, R. 1980. A logic for default reasoning. *Artificial Intelligence* 13(1-2) 81-132.
- [Rusland and for Mosler, 1987] Rusland, Eru. J. and Mosler, Alice G. B. ter, editors 1987. *The Representation of Indefiniteness*. The MIT Press, Cambridge, MA.

- [Reyle, 1993] Reyle, U. 1993. Dealing with ambiguities by underspecification. *Constructions, representation and deduction. Journal of Semantics* 3.
- [Roberts, 1987] Roberts, C. 1987. *Modal Subordination, Anaphora, and Disambiguation*. Ph.D. Dissertation, University of Massachusetts at Amherst, Amherst, MA.
- [Roberts, 1989] Roberts, C. 1989. Modal subordination and pronominal anaphora in discourse. *Linguistics and Philosophy* 12 683-721.
- [Rudinow, 1976] Rudinow, R. 1976. Scope phenomena, "movement transformations," and relative clauses. In Partee, Barbara, editor 1976, *Montague Grammar*. Academic Press 165-176.
- [Rooth, 1987] Rooth, M. 1987. Noun Phrase Interpretation in Montague Grammar. *File Change Semantics, and Situation Semantics*. In Gärdenfors [1987] 237-268.
- [Russell, 1905] Russell, B. 1905. On denoting. *Mind* 14 479-493. Reprinted in *Logic and Knowledge*, ed. R. C. Marsh. London: George Allen and Unwin.
- [Sacerdoti, 1977] Sacerdoti, E. D. 1977. *A Structure for Plans and Behavior*. Elsevier, New York.
- [Safir, 1982] Safir, K. 1982. *Syntactic Chains and the Definiteness Effect*. Ph.D. Dissertation. MIT. Cambridge, MA.
- [Sag, 1976] Sag, Ivan 1976. *Deletion and Logical Form*. Ph.D. Dissertation. MIT, Cambridge, MA.
- [Saurer, 1990] Saurer, W. 1990. A natural deduction system for Discourse Representation Theory. *Journal of Philosophical Logic*.
- [Scha, 1981] Scha, R. 1981. Collective, distributive and cumulative quantification. In Groenendijk, J. A. G., Janssen, T. M. V., and Stokhof, M. B. J., editors 1981, *Formal Methods in the Study of Language*. Walter de Gruyter 483-512.
- [Schank and Abelson, 1977] Schank, R. C. and Abelson, R. 1977. *Scripts, Plans, Goals and Understanding*. Lawrence Erlbaum.
- [Schiffman, 1987] Schiffman, D. 1987. *Discourse Markers*. Cambridge University Press, Cambridge.
- [Schubert and Pelletier, 1982] Schubert, L. K. and Pelletier, F. J. 1982. From English to Logic: Context-free computation of 'conventional' logical translations. *American Journal of Computational Linguistics* 10 165-176.
- [Schubert and Pelletier, 1988] Schubert, L. K. and Pelletier, F. J. 1988. Generically speaking. In Churchus, G., Partee, B. H., and Turner, R., editors 1988, *Properties, Types and Meaning, vol II*. Kluwer, Dordrecht 193-268.
- [Searle, 1969] Searle, J. R. 1969. *Speech Acts*. Cambridge University Press, New York.
- [Shoham, 1988] Shoham, Y. 1988. *Reasoning About Change*. The MIT Press, Cambridge, MA.

- [Small et al., 1988] Small, S. L.; Cottrell, G. W.; and Thachabais, M. K. 1988. *Lexical ambiguity resolution: perspectives from psycholinguistics, neuropsychology, and artificial intelligence*. Morgan Kaufmann, Los Altos, CA.
- [Stecher, 1979] Stecher, R. 1979. Assertion. In Cole, P., editor 1979, *Syntax and Semantics*, volume 9. Academic Press 315-332.
- [Steinberg and Jakobovits, 1971] Steinberg, D. A. and Jakobovits, L. A., editors 1971. *Semantics: An interdisciplinary reader in philosophy, linguistics, anthropology, and psychology*. Cambridge University Press, Cambridge, UK.
- [Stowell, 1981] Stowell, T. 1981. *Origins of Phrase Structure*. Ph.D. Dissertation, MIT, Cambridge, MA.
- [Strawson, 1950] Strawson, P. F. 1950. On referring. *Mind* 59:320-344. Reprinted in *The Philosophy of Language*, ed. A. P. Martinich. New York: Oxford University Press, 1985.
- [Stump, 1985] Stump, Gregory T. 1985. *The Semantic Variability of Absolute Constructions*, volume 25 of *Swedish Language Library*. Dordrecht.
- [Swaney, 1971] Swaney, D. A. 1979. Lexical access during sentence comprehension: reevaluation of context effects. *Journal of Verbal Learning and Verbal Behavior* 18 545-567.
- [Szabolcsi and Zwarts, 1992] Szabolcsi, A. and Zwarts, F. 1992. Weak islands and an algebraic semantics for scope taking. *Natural Language Semantics* 2 1-50.
- [Tachikawa et al., 1979] Tachikawa, M. K.; Stowe, L. A.; and Carlson, G. 1979. Evidence for multiple stages in the processing of ambiguous words in syntactic contexts. *Journal of Verbal Learning and Verbal Behavior* 18(4) 427-440.
- [Thomason, 1974] Thomason, R. H., editor 1974. *Formal Philosophy: Selected Papers of Richard Montague*. Yale University Press, New York.
- [Thompson et al., 1993] Thompson, H.S.; Anderson, A. H.; Bader, M.; Bard, E. G.; Boyle, E.; Dineen-Saddock, G.; Garrod, S.; Isard, S. D.; Korrick, J.; McAllister, J.; Miller, J. E.; Scollin, C.; and Weisert, R. 1993. The HCRC Map Task corpus: A natural spoken dialogue corpus. In *Proceedings International Symposium on Spoken Dialogue*, Waseda University, Tokyo 33-36.
- [Traum and Hinkelman, 1992] Traum, D. E. and Hinkelman, E. A. 1992. Conversation acts in task-oriented spoken dialogue. *Computational Intelligence* 8(3). Special Issue on Non-literal Language.
- [Traum et al., 1994] Traum, D. E.; Allen, J. F.; Ferguson, G.; Heeman, P. A.; Hwang, C. H.; Kato, T.; Martin, M.; Poesio, M.; and Schubert, L. K. 1994. Integrating natural language understanding and plan reasoning in the TRAINS-93 co-ordination system. In *Working Notes AAAI Spring Symposium on Natural Language Understanding in Integrated Systems*, Stanford.

- [Traum, 1993] Traum, D. R. 1993. Mental state in the TRAINS-92 dialogue manager. In *Working Notes AAAI Spring Symposium on Reasoning about Mental States: Formal Theories & Applications*, Stanford.
- [Tyler and Marslen-Wilson, 1982] Tyler, L. K. and Marslen-Wilson, W. D. 1982. Speech comprehension processes. In Mehler, J., Walker, E. C. T., and Garrett, M. F., editors 1982. *Perspectives on Mental Representation: Experimental and Theoretical Studies of Cognitive Processes*. Lawrence Erlbaum, Hillsdale, NJ. 169-184.
- [van Benthem, 1988] Benthem, J. F. A. K. van 1988. *A Manual of Intensional Logic*. CSLI, Stanford, CA, 2nd edition. Distributed by University of Chicago Press.
- [van Doornier, 1991] Doornier, K. van 1991. *On the Composition of Meaning*. Ph.D. Dissertation, University of Amsterdam.
- [van der Sandt, 1988] Sandt, R. A. van der 1988. *Context and Presupposition*. Croom Helm, London.
- [van der Sandt, 1990] Sandt, R. A. van der 1990. Anaphora and accommodation. Workshop on Quantification and Anaphora, Second European Summer School in Language Logic and Information.
- [van Eyck, 1991] Eyck, J. van 1991. Quantification. In Stechow, A. van and Wunderlich, D., editors 1991. *Semantik-An International Handbook of Contemporary Research*. de Gruyter Berlin/New York, chapter 21, 459-487.
- [van Lehn, 1978] Lehn, K. A. van 1978. Determining the scope of English quantifiers. Technical Report AI-TR-483, Artificial Intelligence Laboratory, MIT, Cambridge, MA.
- [Vendler, 1967] Vendler, Z. 1967. *Linguistics in Philosophy*. Cornell University Press, New York.
- [Verkuy and van der Dries, 1991] 'Verkuy, H. J. and Dries, J. van der 1991. The semantics of plural noun phrases. *ILLC preprint*.
- [Verkuy, 1992] Verkuy, H. J. 1992. Some issues in the analysis of multiple quantification with plural NPs. OTS Working Papers OTS-WP TL-92-405, University of Utrecht Research Institute for Language and Speech, The Netherlands. To appear in F. Hamm and F. Hinrichs, editors. *Plural Quantification*, Kluwer.
- [Walker, 1986] Walker, Adrian 1986. Knowledge systems: Principles and practice. *IBM Journal of Research and Development* 30(1), 2-13.
- [Ward and Bumer, 1993] Ward, G. L. and Bumer, B. 1993. *Three sentences and information status*. In *Papers from the Thirtieth annual meeting of the Chicago Linguistic Society*. To appear.
- [Ward, 1985] Ward, G. L. 1985. *The semantics and pragmatics of presupposing*. Ph.D. Dissertation, University of Pennsylvania, Philadelphia.
- [Waterman and Hayes-Roth, 1978] Waterman, D. and Hayes-Roth, F. 1978. *Pattern Directed Inference Systems*. Academic Press, New York.
- [Webber, 1979] Webber, B. L. 1979. *A Formal Approach to Discourse Anaphora*. Garland, New York.
- [Webber, 1987] Webber, B. L. 1987. The interpretation of tense in discourse. In *Proc. ACL-87*, 147-154.
- [Williams, 1988] Williams, W., editor 1988. *Syntax and Semantics 21: Thematic Relations*. Academic Press, New York.
- [Williams, 1977] Williams, E. 1977. Discourse and logical form. *Linguistic Inquiry* 8 101-139.
- [Williams, 1984] Williams, E. 1984. *There-insertion*. *Linguistic Inquiry* 15 131-154.
- [Woods, 1978] Woods, W. A. 1978. Semantics and quantification in natural language question answering. In Yovits, M., editor 1978. *Advances in Computers*. Academic Press, New York 2-64. Also in [Ginsz et al. 1986].
- [Zagawa, 1988] Zagawa, K. 1988. *Verb Phrase Syntax: A Parametric Study of English and Spanish*. Kluwer, Hingham, MA.
- [Zeeval, 1989] Zeeval, H. 1989. A compositional approach to discourse representation theory. *Linguistics and Philosophy* 12:95-131.
- [Zucchi, 1993] Zucchi, A. 1993. The ingredients of definiteness and the definiteness effect. *Natural Language Semantics*. To appear.
- [Zwicky and Sadock, 1975] Zwicky, A. and Sadock, J. 1975. Ambiguity tests and how to fail them. In Kimball, J., editor 1975. *Syntax and Semantics 4*. Academic Press, New York. 1-36.